

NOT FOR CITATION

**IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF CALIFORNIA
SAN JOSE DIVISION**

HEWLETT-PACKARD COMPANY, HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.,

Plaintiffs,

V.

EMC CORPORATION,

Defendant.

EMC CORPORATION.

Counterclaimant.

V.

HEWLETT-PACKARD COMPANY, HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P., COMPAQ COMPUTER CORPORATION,

Counterdefendants.

Case Number C 02-04709 JF

**ORDER CONSTRUING CLAIMS
OF UNITED STATES PATENTS
NO. 4,821,184, NO. 5,237,658, NO.
5,247,618, NO. 5,315,602, NO.
5,325,497, NO. 5,390,327, NO.
5,544,347, NO. 5,659,801, NO.
5,917,253, NO. 6,122,756, NO.
6,269,453, AND NO. 6,356,979; AND
DENYING MOTION TO
PRECLUDE HP FROM
ASSERTING NEW CLAIM
CONSTRUCTION POSITIONS**

On April 5 and 6, 2004, the Court held a hearing for the purpose of construing key terms in the claims of United States Patents No. 4,821,184 (“the ‘184 patent”), No. 658 (“the ‘658 patent”), No. 5,247,618 (“the ‘618 patent”), No. 5,315,602 (“the ‘602

1 patent”), No. 5,325,497 (“the ‘497 patent”), No. 5,390,327 (“the ‘327 patent”), No. 5,544,347
 2 (“the ‘347 patent”), No. 5,659,801 (“the ‘801 patent”), No. 5,917,253 (“the ‘253 patent”), No.
 3 6,122,756 (“the ‘756 patent”), No. 6,269,453 (“the ‘453 patent”), and No. 6,356,979 (“the ‘979
 4 patent”). After consideration of the arguments and evidence presented by the parties and the
 5 relevant portions of the record, the Court construes the disputed terms as set forth below. In light
 6 of its construction of the claims and in order to achieve the most accurate construction possible,
 7 the Court will deny Defendant’s motion to preclude Plaintiffs from asserting new claim
 8 construction positions.¹

10 I. BACKGROUND

11 On September 30, 2002, Plaintiffs and Counterdefendants (“HP”) filed suit against
 12 Defendant and Counterclaimant (“EMC”), alleging infringement of several of the claims of HP’s
 13 ‘658, ‘618, ‘602, ‘253, ‘453, ‘979, and ‘327 patents. EMC counterclaimed alleging that HP
 14 infringed several of the claims of EMC’s ‘184, ‘497, ‘347, ‘801, and ‘756 patents. The ‘184
 15 patent incorporates by reference U.S. Patent No. 4,455,602 (“the ‘5602 patent”), and much of the
 16 Court’s analysis of the ‘184 patent is derived from the ‘5602 patent. The patents generally relate
 17 to devices, systems, and methods for storing or backing up electronic information on memory
 18 systems. The parties seek damages, declaratory judgment, an injunction, treble damages for
 19 willful infringement, attorneys’ fees, and other relief.

20 HP’s asserted patent claims will be described generally first. A general description of
 21 EMC’s asserted claims will follow. These descriptions are not meant to be read as part of the
 22 Court’s claim construction but rather are provided to put the claim construction in context.

23 A. HP’s Asserted Claims

24 The ‘618 patent contains method and apparatus claims for transferring data between

25
 26 ¹ In connection with the claim construction hearing, Defendant also moved to disqualify
 27 Plaintiffs’ expert witness, Dr. Randy Katz. The Court intends to deny the motion and will set
 28 forth its reasoning in a separate order to be issued at a later date. Additionally, because it has not
 considered the declaration of Stephen J. Wallach, the Court declines to address Defendant’s
 objections thereto.

1 storage media. Storage media are materials within a storage device on which data is recorded or
 2 represented (for example, magnetic disks). The purpose of the patented invention is to protect
 3 against loss or corruption of data by maintaining multiple copies of the data. Such multiple
 4 copies are called a “shadow set.” In the main claim, a host processor issues a command for data
 5 on a first storage medium to be copied to a second storage medium. The copies are part of the
 6 shadow set. Whenever one copy is altered, all copies in the shadow set similarly are altered. The
 7 key claim also includes a limitation in which access of the host processor to the data in the
 8 shadow set is “stalled” if the files in the shadow set are in the process of being copied among the
 9 various storage media.

10 The purpose of the invention claimed in the ‘658 patent is to provide an improved method
 11 of connecting several processors (or computers) to several memory storage arrays. The storage
 12 array permits redundant copying of data to several distinct media, thus allowing recovery if one
 13 medium fails. The prior art dealt only with connecting processors to multiple storage devices.
 14 The invention claimed in the ‘658 patent allegedly is novel because processors are connected to
 15 multiple storage *arrays*, each of which has several storage devices for redundant storage. The
 16 change from single storage devices (such as disk drives) to arrays presented a challenge with
 17 respect to the method by which the devices are connected. Accordingly, the inventors claim a
 18 “switching network means” that allows any of several processors to connect to any of several
 19 memory storage devices. Each array has at least one “array controller” through which the array is
 20 connected to the switching network means. The specification describes the switching network
 21 means as an “N x N cross-point switch or an N x N multi-stage switch.” The internal
 22 architecture of the switching network means can be complex.

23 The ‘253 patent focuses on the power supply for the system. Its goal is to provide a
 24 means by which the processor and memory devices retain power even if one of the power
 25 supplies fails, in other words, a redundant physical power supply system. Such a system reduces
 26 the potential for data loss and computer down time. For example, two AC mains may supply
 27 power to two “power supplies,” which then provide DC power to the computer system. If one of
 28 the AC mains fails, the other AC main very rapidly provides power to both of the power supplies

1 “without any substantial loss in transfer of line voltage.” This occurs through an “automatic”
 2 “switching apparatus.” In the preferred embodiment, there are two AC mains and three power
 3 supplies.

4 The ‘602 patent relates to the process of transferring data from a host processor to
 5 redundant memory storage devices. It describes a “redundant array of inexpensive disks”
 6 (“RAID”), on each of which data are stored in a block (“data blocks”). Each disk has a
 7 redundant copy of a data block. The set of redundant copied blocks is called a “stripe.” For each
 8 set of data blocks, a “parity data block” also is generated, and the parity block is part of the
 9 stripe. “Parity data” is information about the copies that enables the system to reconstruct the
 10 data in a data block if the data are lost or corrupted.

11 Associated with the storage devices is a “memory buffer cache,” from which data are
 12 written to the disk drives and which a host computer can access more quickly than it can access
 13 the disk drives. Host “logical I/O” requests (input/output requests) store data to the memory
 14 buffer cache and host “physical I/O” requests store data to the disk drives. Cache storage is more
 15 vulnerable to loss or corruption, while disk drive storage is “non-volatile.”

16 The purpose of the claimed invention is to generate new parity data for a stripe with as
 17 few “I/O requests” as possible whenever there is a change in a data block. Using an I/O request,
 18 the system must read data or parity blocks from the disk drive. For example, if there are five data
 19 blocks in the disk drives with the same copy and the system alters data relating to one of those
 20 blocks, there are two ways to alter the parity data for the stripe. First, each data block may be
 21 read from the disk drives, requiring four I/O requests. Alternatively, the old parity data block and
 22 an old data block may be read from the disk drive and then compared to the new data block, thus
 23 requiring only two I/O requests. If all of the data blocks already are in the memory buffer cache,
 24 then no additional I/O requests would be needed. The claimed invention assesses how to alter a
 25 stripe of blocks using the fewest I/O requests.

26 The ‘453 patent relates to methods for restoring data in a RAID-based memory storage
 27 system if one of the disks fails and for restoring the RAID array itself after the failed disk is
 28 replaced. Here, it is noted that “striping” has the further advantage of speeding up access to data

1 from the memory because the data might be accessed simultaneously from five disks. Data in the
2 stripe are broken into “chunks.” A chunk is a group of consecutively numbered data blocks
3 placed logically or physically consecutively on a single disk. A “block” is the smallest quantity
4 of data that can be read or written to a disk. Each “stripe” consists of chunks written across
5 multiple disks in the array.

6 The RAID-4 and RAID-5 systems include a “parity system.” A “parity block” contains
7 parity information—error correction codes to reconstruct data if a disk fails or a data block
8 otherwise becomes unavailable—for a group of one or more data blocks. A “parity chunk” is
9 one chunk in a strip that contains parity information for the other chunks in the strip. In RAID-4
10 systems, the parity information is stored on a single disk. In RAID-5 systems, it is stored over all
11 of the disks in the array. When a disk fails, it is restored using the redundant data and the parity
12 data.

13 The purpose of the claimed invention is to allow continuous use of the system before the
14 lost or failed disk is restored, that is, while the memory is in a non-redundant state. The system
15 regenerates the lost data from the other data chunks and the parity chunk and then *writes the*
16 *restored data over the parity information.* In this way, all of the data blocks are available, but the
17 memory is no longer redundant (this is the configuration of a RAID-0 array). This process is
18 called “folding.” The patent also discloses a method for restoring the array to its redundant state,
19 or “unfolding.” If the replacement disk contains parity information, that information is
20 determined from the data chunks on the other disks and written to the new disk. If it contains
21 data, then data are read from the disk that originally contained the parity data and written on the
22 replacement disk, and parity data are then determined from the resulting data blocks and stored
23 on the other disks in the appropriate locations.

24 The ‘327 patent is related to the ‘453 patent. Its purpose is to carry out the process just
25 described while allowing the system to remain operational; that is, the array remains accessible to
26 applications. It does this by keeping track of the state of a data block. Specifically, a group of
27 “array state bits” indicate whether a data block that is targeted by an application is in one of the
28 following states: normal, folding, fully folded, or unfolding. The system uses this information to

1 maximize access to the memory array without compromising its integrity.

2 The ‘979 patent describes a system including a memory storage system that is accessible
 3 to multiple host computing systems. The goal of the invention is to permit selective access to
 4 data in the storage system because different host computers and operating systems may conflict
 5 when attempting to access the storage system, thus corrupting the system or limiting
 6 communication ability. Data are organized in the storage system as “logical units” (“LUNs”).
 7 The invention permits certain hosts to access only certain LUNs by using a “configuration table,”
 8 which is stored in memory and which governs the interactions between LUNs and host
 9 computers. Different types of information may be stored in the configuration table.

10 **B. EMC’s Asserted Claims**

11 The ‘801 patent relates to a computer system with peripheral components (such as a
 12 memory storage device, printers, and video devices) that have their own processors and software
 13 programs that control their operation. The goal of the claimed invention is to update the software
 14 in the peripheral components (this software also is known as “microcode”) by delivering it from
 15 the central computer. The process begins when the computer sends an “initiator peripheral
 16 device command” to the peripheral device, which responds by entering a “waiting state,” in
 17 which it is ready to receive the new microcode. The computer then sends a “transfer peripheral
 18 device command,” which contains the new microcode. Finally, the peripheral device replaces the
 19 old microcode with the new.

20 The ‘497 patent discloses a method for labeling a “logical set” of memory. A logical set
 21 is a set of physical disk drives functioning as one logical drive. Each logical set has a “global
 22 identifier” and each member of the set has a “membership signature.” The global identifier and
 23 membership signatures are substantially similar. If one of the disks fails or is taken out of the
 24 memory system and replaced, it no longer would have the same membership signature because it
 25 may no longer correspond to the other members of the logical set. In this way, members of a
 26 logical set are identified and changed members are excluded from the logical set.

27 The ‘184 patent incorporates by reference the hefty ‘5602 patent. It discloses a data
 28 organization method whereby data are organized into “objects,” which are identified by “unique

identifiers” (“UIDs”). “[D]ata is addressed by means of a logical address which specifies the UID of the object containing the data and the offset of the data in the object.” ‘184:3/6-9. In this way, a computer may access data stored in “universal logical memory.”

The '756 patent relates to a self-diagnosing and self-correcting computer system. The computer system automatically tests itself, identifying any flawed components, and then "de-configures" the flawed components without impairing the system as a whole, although the system continues to operate in a degraded state until the component is replaced. A replacement step is not included in the claims.

The ‘347 patent discloses another system for backing up data. A computer with a local primary memory storage system stores data on that primary storage system. The primary storage system has a controller that receives data from the computer and controls their storage. The primary storage controller is connected via high speed communications link to a secondary memory storage controller in a secondary memory storage system that is geographically distinct from the first. The primary storage controller sends a copy of the data to the secondary storage controller, which stores a back-up copy at the geographically distinct location. The secondary storage controller sends an indication to the first controller that it has received the data and that the data are valid.

II. APPLICABLE LAW

Claim construction is a question of law to be decided by the Court. *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 979 (Fed. Cir. 1995), *aff'd*, 517 U.S. 370 (1996). When assessing claim meaning the Court must ask what a person having ordinary skill in the art would understand the claim language to mean at the time of the invention. *DeMarini Sports, Inc. v. Worth, Inc.*, 239 F.3d 1314, 1324 (Fed. Cir. 2001). The Court must look first to the intrinsic evidence of record: the patent claims, the specification, and, if in evidence, the prosecution history. *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996). The general rule is that claim terms are to be given their ordinary and accustomed meaning. *Johnson Worldwide Assocs. v. Zebco Corp.*, 175 F.3d 985, 989 (Fed. Cir. 1999). However, the patentee

1 may choose to be its own lexicographer and may use terms in a manner other than their ordinary
 2 meaning so long as the special definition is stated clearly in the patent specification or file
 3 history. *Vitronics*, 90 F.3d at 1582. The specification acts as a dictionary when it expressly
 4 defines terms used in the claims or when it defines terms by implication, and is considered to be
 5 “the single best guide to the meaning of a disputed term.” *Id.*

6 In most situations, analysis of the intrinsic evidence will resolve any ambiguity regarding
 7 a disputed claim term. *Id.* at 1583. In such circumstances, the Court may not rely on extrinsic
 8 evidence. *Id.* However, where the intrinsic evidence is ambiguous as to a disputed term or the
 9 scope of the invention, the Court may turn to extrinsic evidence such as expert testimony, prior
 10 art, and inventor testimony. *Id.* at 1584. Such evidence may be used to help the Court
 11 understand the claims but may not be used to vary or contradict the claim language. *Id.*

12 Technical treatises and dictionaries may be used at any stage of the claims construction
 13 process. In fact, such sources likely should be used at the initial step of reading the claim
 14 language itself, because “[c]onsulting the written description and prosecution history as a
 15 threshold step in the claim construction process, before any effort is made to discern the ordinary
 16 and customary meanings attributed to the words themselves, invites a violation of [Federal
 17 Circuit] precedent counseling against importing limitations into the claims.” *Texas Digital Sys.,*
Inc. v. Telegenix, Inc., 308 F.3d 1193, 1204 (Fed. Cir. 2003), *cert. denied*, 123 S. Ct. 2230
 18 (2003). In determining the ordinary meaning of a term, the Court may consult technical treatises
 19 and dictionaries publicly available at the time the patent is issued “at any time in order to better
 20 understand the underlying technology and may also rely on dictionary definitions when
 21 construing claim terms, so long as the dictionary definition does not contradict any definition
 22 found in or ascertained by a reading of the patent documents.” *Vitronics*, 90 F.3d at 1584 n. 6;
 23 *see also Teleflex, Inc. v. Ficosa N. Am. Corp.*, 299 F.3d 1313, 1325 (Fed. Cir. 2002). The
 24 intrinsic evidence and dictionaries must be considered together: because “words often have
 25 multiple dictionary definitions, some having no relation to the claimed invention, the intrinsic
 26 record must always be consulted to identify which of the different possible dictionary meanings
 27 of the claim terms in issue is most consistent with the use of the words by the inventor.” *Texas*

1 *Digital Sys., Inc.*, 308 F.3d at 1203.

2 As a general claim construction principle, limitations found only in the written
 3 description of the specification of a patent should not be imported or read into a claim. *In re*
 4 *Donaldson*, 16 F.3d 189, 195 (Fed. Cir. 1994); *Laitram Corp. v. NEC Corp.*, 163 F.3d 1342,
 5 1347 (Fed. Cir. 1998). However, in some instances limitations from the specification may be
 6 imported into a claim. In such instances, the Court “looks to whether the specification refers to a
 7 limitation only as part of less than all possible embodiments or whether the specification read as
 8 a whole suggests that the very character of the invention requires the limitation be a part of every
 9 embodiment.” *Alloc Inc. v. Int'l Trade Comm'n*, 342 F.3d 1361, 1370 (Fed. Cir. 2003). Where
 10 “the specification makes clear at various points that the claimed invention is narrower than the
 11 claim language might imply, it is entirely permissible and proper to limit the claim.” *Id.* A
 12 second exception applies if a claim is expressed in “means plus function” or “step plus function”
 13 format in accordance with 35 U.S.C. § 112, ¶ 6, which states:

14 An element in a claim for a combination may be expressed as a means or a step for
 15 performing a specified function without the recital of structure, material, or acts in
 16 support thereof, and such claim shall be construed to cover the corresponding
 17 structure, material, or acts described in the specification and equivalents thereof.

18 The first step in construing a means-plus-function limitation is to identify the function of
 19 the limitation recited in the claim. *Texas Digital Sys., Inc.*, 308 F.3d at 1208. The next step is to
 20 identify the corresponding structure set forth in the written description necessary to perform that
 21 function. *Id.* “Structure disclosed in the specification is ‘corresponding’ structure only if the
 22 specification or prosecution history clearly links or associates that structure to the function
 23 recited in the claim.” *Id.* (quoting *B. Braun Med., Inc. v. Abbott Labs.*, 124 F.3d 1419, 1424
 24 (Fed. Cir. 1997)). Corresponding structure is limited to that necessary to perform the recited
 25 function and its structural equivalents. *Micro Chemicals, Inc. v. Great Plains Chem. Co.*, 194
 26 F.3d 1250, 1257-58 (Fed. Cir. 1999).

III. DISCUSSION

A. Disputed terms of claim 1 of the '618 patent.

The text of claim 1 is set forth below with the key disputed terms highlighted in bold type.

A method of transferring data between two storage media in a **shadow set** of storage media accessible by one or more **hose [sic] processors** for I/O requests, said method comprising the steps of:

A. receiving a command from one of said **host processors**, said **command specifying data to be transferred** from a first said storage medium to a second said storage medium;

B. transferring said data specified in said command from said first storage medium to said second storage medium in a series of subtransfers, each of said subtransfers transferring a portion of said specified data; and

C. processing one or more I/O requests received from said **host processors** for access to said **shadow set** by, for each received I/O request:

a. implementing said I/O request if said I/O request does not involve a section of said **shadow set** currently involved in one of said subtransfers; and

b. stalling said I/O request if said I/O request involves a section of said **shadow set** currently involved in one of said subtransfers, and implementing said I/O request when said one of said subtransfers has been completed.

1. “host processor”

HP asks the Court to construe this term as “a processor that controls or accesses all or part of a user application network.” EMC proposes: “a processor in a host computer that can issue read or write requests to the first and second storage media.” HP argues that “the [parties’] only disagreement concerns whether or not the processor must reside in something called a “host computer.” Plaintiffs’ and Counterclaim Defendants’ Corrected Opening Claim Construction Brief on Key Terms (“HP’s Opening Brief”), p. 6 n.4. As defined by the claim, a host processor is a processor that is capable of sending commands to a storage medium or requests for access that are received by a shadow set. HP’s proposal adds a reference to a “user application network.” EMC’s proposed language adds the requirement that the host processor reside in a host computer.

According to HP, the dictionary meaning of “host processor” is “a processor that controls all or part of a user application network.” HP’s Opening Brief, p. 6; Declaration of Christopher D. Landgraff (“Landgraff Decl.”), Separate Appendix, vol. 1, p. 0195. However, the claim language does not include reference to a “user application network,” and it is not clear what a

1 “user application network” for the purposes of the claim would be, although it could be the
 2 claimed system as a whole. Instead, the claim language refers to commanding a storage medium
 3 and accessing a shadow set. Similarly, while the written description includes diagrams
 4 suggesting that a host processor may be part of a “host,” the claim language does not include
 5 reference to a “host computer.” The claim language simply has no location limitation, and while
 6 the word “host” may limit “processor,” it need not do so by geography. Looking at all of the
 7 dictionary definitions provided by the parties, it is clear that a “host processor” is part of a
 8 network or system and that it controls or accesses other parts of the system. *See id.* The plain
 9 language of the claim indicates that “host processor” refers to a specific type of processor, one
 10 that sends commands to the memory aspects of the system. The claim language does not specify
 11 the physical location of the host processor.

12 Accordingly, the appropriate construction of the “host processor” as used in claim 1 is “a
 13 processor that can send a command to a storage medium or requests for access that are received
 14 by a shadow set.”

15 **2. “command specifying data to be transferred”**

16 HP asks the Court to construe this term as “an order identifying data to be transmitted
 17 from a first storage medium to a second storage medium.” EMC proposes: “command
 18 specifying the data to be transferred by reference to the starting block numbers on the source and
 19 target and the quantity of data to be copied.”

20 According to the plain language of the claim, a “command specifying data to be
 21 transferred” is a command sent from a host processor that specifies which data are transferred
 22 from a first storage medium to a second storage medium. HP’s proposed construction thus
 23 adheres generally to the plain language of the claim. EMC, on the other hand, seeks to import
 24 limitations from the specification. HP argues that to “specify” is to “identify,” while EMC seeks
 25 to include in the definition the method by which “specifying” occurs, thus not only defining the
 26 term but also importing a method limitation. The parties agree that at a minimum “specifying”
 27 means “identifying, naming, or stating explicitly or in detail.” Joint Claim Construction and
 28 Prehearing Statement (“JCCPS”), Ex. B. This is consistent with the plain language of the claim

1 and is sufficient for purposes of claim construction. There is no ambiguity that requires reference
 2 to a particular method or embodiment for identifying explicitly or in detail. Nor is there any
 3 apparent reason for exchanging the word “transferred” for “transmitted” and “order” for
 4 “command” as proposed by HP.

5 Accordingly, the appropriate construction of “command specifying data to be transferred”
 6 as used in claim 1 is “a command sent from a host processor that identifies, names, or states
 7 explicitly or in detail which data is transferred from a first storage medium to a second storage
 8 medium.”

9 **3. “shadow set”**

10 HP asks the Court to construe this term as: “Two or more copies of data that may be
 11 located in different places. If a change is made to any copy, that change is reflected in the
 12 others.” EMC proposes: “Two or more identical copies of data that has been written to a storage
 13 medium, and that are treated as one from the host processor’s point of view. If a change is made
 14 to any copy, that change is automatically made to the other copies.”

15 At first glance, the plain language of the claim appears to indicate that a “shadow set” is
 16 composed of *storage media* that are accessible by one or more host processors for I/O requests.
 17 The claim states that data are transferred between different storage media within a shadow set
 18 and that the storage media are accessed by a host processor. The shadow set thus would
 19 comprise defined storage media that contain the data to be transferred within the defined storage
 20 media. However, in the context of the claim, the language “shadow set of storage media” could
 21 mean a limited set consisting of part of the storage media that contain a shadow set. In other
 22 words, a shadow set could be composed of *data* that are moved between storage media, in which
 23 case the identity of the storage media would be irrelevant. Indeed, the parties characterize a
 24 shadow set as “copies of data” either written to a storage medium or located in different places.
 25 Thus, it is unclear from the claim language itself whether the shadow set is composed of storage
 26 media or data.

27 The specification indicates that “shadow set” refers to data. For example, the “preferred
 28 embodiment is described in connection with a system for establishing and maintaining one or

1 more duplicate or ‘shadow’ copies of stored data.” ‘618:1/27-30. Similarly, “systems have been
 2 developed that create multiple copies of stored data, usually on separate storage devices
 3 Such multiple copies are known as the shadow set.” ‘618:1/42-47. However, the parties each
 4 seek to add an additional limitation relating to whether all copies in a shadow set are identical or
 5 reflect each other and whether a change to one copy results in a change on all copies. The
 6 specification provides that “[i]n a shadow set, typically data that is stored in particular blocks on
 7 one member of the shadow set is the same as data stored in corresponding blocks on the other
 8 members of the shadow set.” ‘618:1/47-50. Thus, because the shadow set by definition contains
 9 redundant copies of the data, each copy must be the same, and necessarily, any change in one
 10 copy correspondingly is made in other copies within the shadow set. If a change in one copy is
 11 not made to the others, then that copy no longer can be a member of the shadow set. Dictionary
 12 definitions support this interpretation. For example, “shadow” is partially defined as “a link
 13 between duplicate objects. . . . If a change is made in either the duplicate or the original, the
 14 change takes effect in the other as well.” Landgraff Decl., Separate Appendix, vol. 1, p. 0197
 15 (IBM Dictionary of Computing).

16 EMC proposes adding additional limitations, requiring the copies of data to be “treated as
 17 one from the host processor’s point of view” and updated “automatically.” However, nothing in
 18 the claim language or specification supports or requires limiting the claim in this way.

19 Accordingly, a person of ordinary skill in the art would have understood the “shadow set”
 20 term of claim 1 to mean: “Copies of data located on one or more storage media. If a change is
 21 made to any copy, that change is reflected in the others.”

22 **B. Disputed terms of claims 1 and 14 of the ‘658 patent.**

23 **1. “switching network . . . for establishing a communications link”**

24 The text of claim 14 is set forth below with the key disputed term highlighted in bold
 25 type.

26 A multiprocessor computer system comprising:
 27 a. a plurality of processing units;
 28 b. at least one data storage array system, each having at least one array
 controller;
 29 c. at least one **switching network**, coupled to the plurality of processing

1 units and to at least one array controller of at least one data storage array system,
 2 **for establishing a communications link** between at least one selected processing
 3 unit and at least one data storage array system.

4 HP asks the Court to construe this term as: “A component or set of components to which
 5 multiple nodes can attach and that selectively enables giving, receiving, or exchanging of
 6 information, signals and/or data among attached nodes.” EMC proposes: “A circuit switch, not
 7 contained within a data storage array system, that is capable of establishing a dedicated path from
 8 any input node of the switch to any output node of the switch.” The critical dispute between the
 9 parties is whether the term “broadly covers both ‘circuit switching’ and ‘packet switching’” or
 10 whether it is limited to “circuit switching.” EMC’s Response to HP’s Opening Claim
 Construction Brief (“EMC’s Response”), p. 12.

11 The plain language of the claim indicates that a switching network establishes a
 12 communications link between the processing units and the data storage array system by coupling
 13 to them and that a plurality of processing units is coupled to a switching “network.” Thus, there
 14 must be multiple connection points between the processing units and the switching network.
 15 This conclusion is consistent with HP’s construction.

16 Nothing in the claim language limits “switching network” to a particular type of “switch.”
 17 The switching network only need be coupled to the processing units and the data storage array
 18 system. Thus, both circuit and packet switching could occur within the switching network. In
 19 fact, the specification provides support for multiple types of switching. For example, the
 20 specification states that it does not limit the type of switching to the disclosed embodiments,
 21 ‘658:4/44-47, Figure 2 appears to depict at least direct circuit switching, and Figure 3B depicts
 22 “multi-stage switch architecture,” which consists of a matrix of “selectors” through which signals
 23 are transferred. In Figure 3B, each signal may travel through multiple selectors to its final
 24 destination.

25 Although the specification does not support importation of the limitation suggested by
 26 EMC, it does provide clarification as to the intended meaning of “switching network.” That is, a
 27 switching network must be capable of coupling any node to any other node. *See, e.g.*, ‘658:5/25-
 28 26 & 5/38-40. Dictionary definitions are consistent with this view. For example, “switch” is

1 defined as “a network infrastructure component to which multiple nodes attach,” and “network”
 2 is defined as “an interconnect that enables communication among a collection of attached
 3 nodes.” *See Landgraff Decl., Separate Appendix, vol. 2, pp. 0324-26 (A Dictionary of Storage*
 4 *Networking Terminology).* A person of ordinary skill in the art would understand a switching
 5 network to mean a component or a set of components to which multiple nodes can attach and that
 6 selectively enables transferring of information, signals, or data among the attached nodes.

7 EMC also contends that the switching network must not be located within the data
 8 storage array system. However, nothing in the intrinsic evidence supports such a limitation. In
 9 the prosecution history cited by EMC, the patentee distinguished the invention from the prior art
 10 on the ground that a switching network was used between the processors and the disk storage
 11 array system, while the prior art described the use of a switch between the array controller and
 12 the storage units. *See Landgraff Decl., Separate Appendix, vol. 2, p. 0285.* The distinction does
 13 not concern the *physical* location of the switch but rather where the switch is placed *logically* in
 14 the system in relation to each component. Logically, the switching network must be between the
 15 processors and the disk storage array system.

16 Accordingly, the Court construes this term as “a component or set of components to
 17 which multiple nodes can attach and that selectively enables transferring of information among
 18 attached nodes.”

19 **2. “switching network means . . . for establishing a communications link”**

20 The text of claim 1 is set forth below, with the key disputed terms highlighted in bold
 21 type.

22 A multiprocessor computer system comprising:
 23 a. a plurality of processing units;
 24 b. at least one data storage array system, each having at least one array
 25 controller;
 26 c. **switching network means**, coupled to the plurality of processing units
 27 and to at least one array controller of at least one data storage array system, **for**
 28 **establishing a communications link** between at least one selected processing
 unit and at least one data storage array system.

29 This claim is expressed in means plus function format, and the Court therefore construes
 30 it pursuant to 35 U.S.C. § 112, ¶ 6. HP defines the function as “establishing a communications

1 link between at least one selected processing unit and at least one data storage array system.”

2 Because HP’s proposal is consistent with the plain language of the claim, the Court will adopt it.

3 HP defines the corresponding structures as found in “Figs. 2, 3A, 3B, 4; Col. 5:1-6:45.”

4 *See* HP’s Opening Brief, p. 13. HP thus asks the Court to construe this term as: “A component
5 or set of components to which multiple nodes can attach and that selectively enables giving,
6 receiving, or exchanging of information, signals and/or data among attached nodes, having
7 structures that are the same as or substantially equivalent to those found in Figures 2-4 and
8 described at Columns 5:1-6:45.” EMC defines the structure more narrowly, proposing: “A
9 circuit switch, not contained within a data storage array system, that is either a cross-point or
10 multi-stage switch that is capable of establishing a dedicated path from any input node of the
11 switch to any output node of the switch, as described in Col. 5, ll 4-51 and shown in Figs. 3A and
12 3B.”

13 EMC’s argument is persuasive in part. The structures described in Figures 3A and 3B
14 and in Column 5, lines 4 through 67 clearly are associated with the function of establishing a
15 communications link between at least one selected processing unit and at least one data storage
16 array system. *See Texas Digital Sys., Inc.*, 308 F.3d at 1208. In contrast, while Figures 2 and 4
17 show a “switching network,” they do not provide the structure of the “switching network,” but
18 instead show the “switching network” as part of the overall system configuration. Similarly, the
19 text in the written description corresponding to those figures does not clearly link the function of
20 the switching network means to its structure.

21 When the patent specification describes multiple embodiments, each of those
22 embodiments may be claimed pursuant to 35 U.S.C. § 112, ¶ 6. *Serrano v. Telular Corp.*, 111
23 F.3d 1578, 1583 (Fed. Cir. 1997). The structures described in Figures 3A and 3B and in Column
24 5, lines 4 through 67 generally “may be of any suitable N x N type, capable of directly coupling
25 any node to any other node.” The figures and written description also describe specifically “an N
26 x N cross-point switch or an N x N multi-stage switch” as disclosed in Column 5, lines 7 through
27 51. The written description additionally indicates that the switching network may comprise fiber
28 optic or wired links and is fault-tolerant and that the switch may be a dual switching network.

1 Accordingly, the structure of the “switching network means for establishing a communications
 2 link” is “any suitable N x N type switching network, capable of directly coupling any node to any
 3 other node in addition to the specific embodiments disclosed in Figures 3A and 3B and Column
 4 5, lines 7 through 67.”²

5 **C. Disputed terms of claims 1 and 12 of the ‘253 patent.**

6 The text of claim 1 is set forth below with the key disputed terms highlighted in bold
 7 type.

8 A circuit comprising:

- 9 (a) first and second energized AC feed circuits;
- 10 (b) **first and second receiving circuits connected to and energized by,**
respectively, the first and second feed circuits; and,
- 11 (c) first switching apparatus which automatically connects the second feed
 circuit to the first receiving circuit in the event the first feed circuit is de-
 energized, whereby both the first and second receiving circuits are energized by
 the second feed circuit without any **substantial loss in transfer of line voltage.**

12 The text of Claim 12 is set forth below, with the key disputed terms highlighted in bold
 13 type.

14 A redundant electrical system comprising:

- 15 (a) first and second energized AC energy sources;
- 16 (b) first and second energy transfer circuits connected to and energized by,
 respectively, the first and second energy sources; and,
- 17 (c) a first relay connected to and energized by the first energy source for
 automatically disconnecting the first energy source from the first transfer circuit
 and automatically connecting the second energy source to the first transfer circuit
 in the event the first energy source is de-energized, whereby both the first and
 second transfer circuits are energized by the second energy source without any
substantial loss in the transfer of line voltage.

20 1. **“first and second receiving circuits connected to and energized by,**
 21 **respectively, the first and second feed circuits”**

22 HP asks the Court to construe this term as follows: “The first receiving circuit is
 23 connected to and receives electricity or is currently capable of receiving electricity from the first
 24 feed circuit and the second receiving circuit is connected to and receives electricity or is currently
 25 capable of receiving electricity from the second feed circuit.” EMC proposes: “The first

27 ² Because it need not determine in the context of claim construction whether Figure 3B (a
 28 multi-switch) includes the possibility of “packet switching,” the Court will reserve judgment on
 the issue. The parties may describe the disclosed structure further in the infringement analysis.

1 receiving circuit is physically connected to and is receiving electricity from the first AC feed
 2 circuit. At the same time, the second receiving circuit is physically connected to and is receiving
 3 electricity from the second AC feed circuit.”

4 Essentially, the parties dispute the meaning of the phrase “energized by.” HP asks the
 5 Court to construe “energized by” as “receives electricity or is currently capable of receiving
 6 electricity.” EMC argues that the correct construction is “is receiving electricity.” Under EMC’s
 7 proposal, claim 1 would not protect an invention in which the receiving circuit is not currently
 8 receiving electricity from its respective AC feed circuit.

9 The meaning of “energized by” in the claim language is ambiguous. The Court thus turns
 10 to the specification for clarification. The written description indicates that the “AC mains [are]
 11 feed circuits which provide voltage,” ‘253:3/9-10, and that “when both AC mains are energized,
 12 all three power supplies receive voltage,” ‘253:3/34-36. Thus, for the purposes of the claim, a
 13 person of ordinary skill in the art would understand “energize” to mean “to provide voltage to.”
 14 “Energized by” thus generally means “receives voltage from.”

15 The claim states that a receiving circuit is “connected to and energized by” a feed circuit.
 16 Accordingly, a receiving circuit is “connected to and receives voltage from a feed circuit.”
 17 Nothing in the claim language indicates that a receiving circuit must *constantly* receive voltage
 18 from the same feed circuit. Indeed, the claim language as a whole describes a situation in which
 19 both receiving circuits are energized by the second feed circuit “in the event the first feed circuit
 20 is de-energized.” The specification also suggests that a feed circuit might not always provide
 21 voltage to a receiving circuit. For example, an AC main might fail, in which case it could not
 22 provide voltage to a receiving circuit. *See, e.g.*, ‘253:3/29, 4/6-14, & 4/34-45. While the
 23 prosecution history shows that the patent examiner found the term “capable of being energized”
 24 indefinite, the meaning of “energize” in the context of the claim and the specification must
 25 include scenarios in which the receiving circuit is not always energized by the same feed circuit.
 26 There is no evidence in the prosecution history that the patentees stated that a certain feed circuit
 27 must continually energize a receiving circuit. Rather, the evidence suggests that they rewrote the
 28 claims to clarify that “the receiving circuits of claim 1 . . . are connected to and energized by the

1 respective feed circuits.” See Landgraff Decl., Separate Appendix, vol. 3, p. 0386. In other
 2 words, they did not intend to claim continual energizing, but rather to clarify which circuit
 3 energizes the receiving circuit for the purposes of subparagraph (b) of claim 1.

4 EMC also contends that, based on the claim language and the prosecution history, the two
 5 receiving circuits must be energized “at the same time.” While it is true that the patentees argued
 6 that their “invention keeps at least two load circuits energized at all times, regardless of whether
 7 one energy source fails or not,” *id.* at 0388, nothing in the disputed claim language requires that
 8 the two be energized at the same time by distinct AC feed circuits. The claim language merely
 9 indicates that the two receiving circuits are connected to and energized by feed circuits. As
 10 noted, the receiving circuits may not always be energized by their respective feed circuits.
 11 Although EMC’s argument initially appears consistent with subsection (b) of the claim and the
 12 relevant prosecution history, it nonetheless does not make sense to require the two receiving
 13 circuits to be energized at the same time by *respective* feed circuits because, if a receiving circuit
 14 ceases receiving voltage from a feed circuit, “automatic switching capability” permits the
 15 invention to keep at least two receiving circuits energized “at all times” through a *different*,
 16 functioning feed circuit. *Id.*

17 The Court thus construes this term as: “The first receiving circuit is connected to and
 18 receives electricity or is capable of receiving electricity from the first feed circuit and the second
 19 receiving circuit is connected to and receives electricity or is capable of receiving electricity from
 20 the second feed circuit.”

21 **2. “substantial loss in transfer of line voltage”**

22 HP asks the Court to construe this term as: “The voltage at the destination of the transfer
 23 is substantially less than that at the source.” With respect to claim 1, EMC proposes: “A
 24 substantial loss in AC line voltage when switching from the first AC feed circuit to the second
 25 AC feed circuit. A loss in AC line voltage of 3dB or more is a substantial loss.” With respect to
 26 claim 12, EMC proposes: “A substantial loss in AC line voltage when switching from the first
 27 AC energy source to the second AC energy source. A loss in AC line voltage of 3dB or more is a
 28 substantial loss.”

1 Subparagraphs (c) of claims 1 and 12 indicate that the receiving circuits of claim 1 and
 2 the transfer circuits of claim 12 are energized without substantial loss from a single feed circuit
 3 or energy source, respectively, when compared to the amount of voltage transfer received from
 4 both feed circuits or energy sources. That is, each receiving or transfer circuit must be energized
 5 in substantially the same amount after failure of one of the two feed circuits or energy sources.
 6 Accordingly, for the purposes of this term, the voltage signal at the final destination (either the
 7 receiving circuits or the transfer circuits) prior to the de-energization of the first feed circuit or
 8 first energy source is compared to the voltage signal at the same final destination after the de-
 9 energization of the first feed circuit or first energy source.

10 The parties dispute whether there may be a “substantial loss in the transfer of line
 11 voltage” *during* the switch. Clearly, the object of the claimed invention is to prevent loss of
 12 energy to a circuit if an AC feed line or energy source fails. The specification describes an
 13 example of a switch that can achieve this object: for example, the “switching speed of relay 65
 14 must be less than about one signal cycle of the AC main (i.e., about 20 ms), so that continuity of
 15 the signal is maintained for the power supply after switching AC mains.” ‘253:3/60-63. Thus,
 16 consistent with the claim language, the specification describes a device in which transfer of line
 17 voltage signal is never substantially lost, either during or after the switch, as measured against the
 18 extent to which the circuit was energized prior to the switch. Although there may be energy loss
 19 during the switch, that loss is not substantially less with respect to the final destination—because
 20 of the nature of the AC signal—than what would have occurred if there had been no switch.
 21 Although the Court concludes that the claim language indicates that substantial loss is measured
 22 after “the first feed circuit [or energy source] is de-energized” rather than after the switch is
 23 made, it must be noted that there may be energy loss associated with the switch as long as such
 24 loss occurs during a specific portion of the AC cycle so that there is not substantial loss of
 25 voltage at the final destination when compared to the voltage received prior to the de-
 26 energization of the AC energy source.

27 EMC also seeks to add a limitation defining “substantial loss” as a loss in AC line voltage
 28 of 3 dB or more. The term “substantial” is ambiguous. The specification does not expressly

1 define—indeed it does not discuss at all—the concept of “substantial loss.” In this context, the
 2 written description instead consistently describes the invention as providing for *redundant*
 3 sources of energy should a source fail. For example, various embodiments of the invention are
 4 described, in which “power supplies . . . remain energized for supplying power to the system,”
 5 ‘253:4/4-5, the circuit “always suppl[ies] AC current to the off (third) power supply . . . ,”
 6 regardless of a loss of either of the redundant AC mains,” ‘253:4/12-14, “no matter which AC
 7 main fails, all three power supplies remain powered for system 60 via the other live redundant
 8 main,” ‘253:4/32-34, if a first AC main fails, a second AC main “provid[es] continued voltage,”
 9 ‘253:4/43, and the system may “retain energized both power supplies . . . in the event of failure
 10 of either AC main,” ‘253:4/48-49.

11 While the written description never quantifies the energy loss resulting from the use of
 12 “redundant” energy supplies, the patentee discloses an invention in which the circuits may
 13 function normally as if a source of energy had not been lost. Thus, it is clear that, whatever
 14 precise quantity the parties wish to assign to the term “substantial loss in the transfer of line
 15 voltage,” it must have significance with respect to the functioning of the device. That is, for
 16 “substantial loss” to occur, the output of the electrical system or functioning of whatever device
 17 is powered by the claimed invention must be impaired by loss of the redundant energy source.
 18 Examples of such impairment are described in the specification. For example, loss of energy
 19 could result in “loss of data and down time upon failure.” ‘253:1/14-15. Accordingly, in the
 20 claimed invention, the output of the electrical system is not affected by a failed AC line because,
 21 through the use of redundancies, there is not a substantial loss in the transfer of line voltage.

22 EMC argues that substantial loss must be quantified as “3 dB” because a reference
 23 discussing a signal loss of up to 3 dB was presented to the patent examiner during prosecution.
 24 However, this reference merely showed that a 3 dB signal loss is substantial, a characterization
 25 with which HP agrees. Plaintiffs’ and Counterclaim Defendants’ Reply Brief on Key Terms, p.
 26 19. It did not preclude an argument that something less than 3 dB also is substantial. EMC does
 27 not explain what “3 dB” means in the context of the present invention. Because the reference to
 28 the prior art in the prosecution history did not address the issue of quantifying “substantial loss”

1 and because the meaning of “3 dB” is ambiguous, the Court finds the prosecution history
 2 unhelpful in providing additional clarification.

3 Accordingly, the Court construes this term as: “Loss of receipt of energy—as measured
 4 by the voltage at the final destination (either the receiving circuits or the transfer circuits) after
 5 the de-energization of the first feed circuit or first energy source compared to the voltage at the
 6 same final destination prior the de-energization of the first feed circuit or first energy
 7 source—sufficient to impair the proper functioning of whatever device is powered by the claimed
 8 invention during and after loss of the redundant energy source.”

9 **D. Disputed terms of claim 1 of the ‘602 patent.**

10 The text of claim 1 is set forth below with the key disputed terms highlighted in bold
 11 type.

12 An apparatus for generating parity data for a plurality of disk drives, each of the
 13 plurality of disk drives divided into a plurality of blocks, the plurality of blocks
 14 further organized into a plurality of stripes, each of the plurality of stripes
 15 including one block from each of the plurality of disk drives, one of the blocks of
 16 each stripe storing parity data generated from the data stored in the other blocks of
 17 the stripe, the apparatus comprising;

18 **memory means** for storing data, said memory means partitioned into
 19 blocks compatible with the block structure of the plurality of disk
 20 drives;

21 means for identifying an updated memory block storing data to be written
 22 to one of the plurality of disk drives;

23 means for identifying a predetermined stripe, said predetermined stripe
 24 including a disk block to be overwritten by said updated memory block;

25 **first means for determining the number of I/O requests required to
 26 aggregate in said memory means all of the data, except the parity
 27 data, of said predetermined stripe;**

28 **second means for determining the number of I/O requests required to
 29 aggregate in said memory means said block to be overwritten and the
 30 parity block of said predetermined stripe; and**

31 means for choosing said first means OR said second means to perform said
 32 I/O requests by choosing the means which requires the fewest number of
 33 I/O requests to aggregate in said memory means the blocks of said
 34 identified stripe necessary to generate the parity data of said predetermined
 35 stripe.

1 1. **“memory means”**

2 Because this term is expressed in means plus function format, the Court must construe it
 3 pursuant to 35 U.S.C. § 112, ¶ 6. HP characterizes the function as “memory” and asks the Court
 4 to construe the term as: “A unit of a computer that preserves data for retrieval having structures
 5 that are the same as or substantially equivalent to those found in Figure 1 at 11, 40 and described
 6 at Col. 2:44-52.”³ EMC contends that the function is “storing data” and the corresponding
 7 structure is “cache (40) connected to a host computer (10) through a system bus (30), and located
 8 between the host computer (10) and the disk device (20).”

9 The function of the “memory means” clearly is data storage. The corresponding structure
 10 in the specification that is associated or linked with this function is described in Figure 1 and in
 11 Column 2, lines 44 through 52. The memory means that is partitioned into blocks is shown in
 12 Figure 1 as structure (40) and also is known as the “cache.”

13 While the parties agree that the underlying structure is the cache (40), *see, e.g.*, Transcript
 14 of April 6, 2004 Hearing, 328:5-6, they disagree as to whether the structure has a location
 15 limitation. EMC argues that the block diagram shows that the cache must be located between the
 16 host computer and the disk device and that it must be connected through a bus. However,
 17 nothing in the claim or specification supports this argument. Figure 1 is not a drawing of a
 18 likeness of the system as it would exist in three-dimensional space; rather, it is a schematic
 19 representation that shows how the components are connected *logically*. Figure 1 equally might
 20 indicate that the disk device is located between the host computer and the cache. Nor does the
 21 written description contain an unambiguous indication of how the components are located
 22 physically with respect to each other. For example, it states that the host computer is connected
 23 to the disk device via a communication bus, ‘602:2/49-50, and then states that the computer
 24 system includes a cache that also is connected to the system bus, ‘602:2/51-52. Accordingly, it
 25 would be inappropriate for the Court to include a location limitation in the identification of the
 26 underlying structure.

27
 28 ³ HP acknowledges that it erred in referring to number 11 and that number 40 is the
 correct structure. *See* Transcript of April 6, 2004 Hearing, 315:20-316:3.

1 The Court therefore construes this term as having the function of “data storage.” The
 2 corresponding structure is the “cache” as shown in Figure 1 and described in Column 2, lines 51-
 3 52.

4 **2. “first means for determining the number of I/O requests required to
 5 aggregate in said memory means all of the data, except the parity data, of
 6 said predetermined stripe”**

7 This term also is expressed in means plus function format, and the Court thus must
 8 construe it pursuant to 35 U.S.C. § 112, ¶ 6. The function of the “first means” is “determining
 9 the number of I/O requests required to aggregate in the memory means all of the data, except the
 10 parity data, of the predetermined stripe.”

11 HP contends that the corresponding structures: “are the same as or substantially
 12 equivalent to those found in Figure 4 and described at Col. 4:21-26; Col. 5:10-15; 23-26; 34-52.”
 13 EMC argues that the corresponding structure is “a host computer (10) programmed to perform
 14 the algorithm described at column 5, lines 34-52 and illustrated in figure 4.” The parties dispute
 15 only whether the computer *system* generally or the “host” component of the computer system
 16 alone performs the function. While the direct references to Figure 4 in the specification discuss
 17 the “computer” or the “system,” and not the “host,” *see* ‘602:5/10, 5/36, & 5/42, the written
 18 description does not state what part of the system performs the function. Elsewhere, however,
 19 the written description states that the “computer system 1 includes a central processor unit or
 20 ‘host’ 10 having primary temporary data storage, such as memory 11,”⁴ ‘602:2/46-48; *see also*
 21 ‘602:Figure 1, and that the “host” is the part of the system that issues I/O requests, ‘602:2/58.

22 Moreover, identifying the structure as the system rather than the host standing alone may
 23 lead to invalidation of the claim on grounds of indefiniteness because the patent does not
 24 describe how the “computer system,” as opposed to the “host,” would achieve the stated
 25 function. Controlling Federal Circuit authority holds that:

27 ⁴ The host clearly includes more than a “central processor unit.” For example, it includes
 28 memory 11. However, the host does not include the cache 40 or disk drive 20. *See* ‘602:Figure
 1.

1 The definiteness requirement is to ensure that the claims are written in such a way
 2 that they give notice to the public of the extent of the legal protection afforded by
 3 the patent, so that interested members of the public, e.g., competitors of the patent
 4 owner, can determine whether or not they infringe.

5 *All Dental Prodx, LLC v. Advantage Dental Prods., Inc.*, 309 F.3d 774, 779-80 (Fed. Cir. 2002).

6 In order to avoid a finding of indefiniteness, the Court must thoroughly “attempt to understand
 7 the meaning of a claim [in order] to resolve material ambiguities.” *Id.* at 780. Because the claim
 8 at issue is a means plus function claim, the Court must look to the specification for clearly
 9 associated structure. The specification clearly describes only a system in which the “system”
 10 includes a host that issues the I/O requests. Accordingly, the corresponding structure is “the host
 11 in the computer system 1 programmed to perform the algorithm described in Column 5, lines 10
 12 through 15, 23 through 26, and 34 through 52 and illustrated in Figure 4.”

13 **3. “second means determining the number of I/O requests required to
 14 aggregate in said memory means said block to be overwritten and the parity
 15 block of said predetermined stripe”**

16 Once again, this term is expressed in means plus function format, and the Court construes
 17 it pursuant to 35 U.S.C. § 112, ¶ 6. The Court construes the function of the “second means” as
 18 “determining the number of I/O requests required to aggregate in the memory means said block
 19 to be overwritten and the parity block of the predetermined stripe.”

20 HP contends that the corresponding structures: “are the same as or substantially
 21 equivalent to those found in Figure 5 and described at Col. 4:27-36; Col. 5:15-19; 26-30; 5:53-
 22 6:18.” EMC argues that the corresponding structure is “a host computer (1) programmed to
 23 perform the algorithm described at column 5, line 53 through column 6, line 18 and illustrated in
 24 figure 5.” The parties again dispute whether it is the computer system generally or solely the
 25 “host” component of the computer system that performs the function.

26 As discussed above, while it is true that the direct references to Figure 5 in the
 27 specification discuss the “computer” or the “system,” and not the “host,” *see* ‘602:5/57, 5/60,
 28 5/63, 7/1, & 7/7, the written description states that the “computer system 1 includes a central
 processor unit or ‘host’ 10 having primary temporary data storage, such as memory 11,”

1 ‘602:2/46-48; *see also* ‘602:Figure 1, and that the “host” is the part of the system that issues I/O
 2 requests, ‘602:2/58. Moreover, as noted above, identifying the structure as the system rather than
 3 the host could result in a determination that the claim is indefinite.

4 Accordingly the Court construes the corresponding structure as “the host in the computer
 5 system 1 programmed to perform the algorithm described in Column 5, lines 16 through 19,
 6 Column 5, lines 26 through 30, and Column 5, line 53 through Column 6, line 18 and illustrated
 7 in Figure 5.”

8 **4. “I/O request”**

9 HP asks the Court to construe this term as: “A directive to any operation, program, or
 10 device that transfers data to or from a computer.” EMC proposes: “A request issued by a host
 11 computer, such as a read or write request, to transfer data.”

12 Although the claim indicates that I/O requests are used to aggregate data in the memory
 13 means, the meaning of the term itself is not clear from the claim language alone. The
 14 specification clarifies that I/O requests are used both to write new data to the cache and disk
 15 device and to organize and distribute data to those structures. For example, “the host 10 issues
 16 I/O requests, such as reads and writes, to transfer data between memory 11 and the disk device 20
 17 via the bus 30 and cache 40.” ‘602:2/58-60. “Host 10 ‘logical’ I/O write requests store the new
 18 data in the cache 40, and the ‘physical’ I/O write requests transfer the data from the cache 40 to
 19 the disk drives.” ‘602:3/6-8. “With striping, a host I/O request distributes the data to be
 20 transferred across the disk drives.” ‘602:3/62-64. The subsequent portions of the written
 21 description describe how the computer system determines and attempts to minimize the number
 22 of required I/O requests. Thus, the specification indicates that I/O requests are issued from the
 23 host 10, and the computer system attempts to minimize the number of required requests.

24 A location limitation is not required for the purpose of construing the term. This is not a
 25 means plus function term, and the claim language does not include the location from which I/O
 26 requests are issued. Instead, the claim is directed toward organizing the memory of the system,
 27 and it does not identify the particular computer components that direct such organization.
 28 Definitions from dictionaries support this view. “I/O” is defined as the

1 process of moving data between a computer system's main memory and an
 2 external device or interface such as a storage device, display, printer, or network
 3 connected to other computer systems. I/O is a collective terms for *reading*, or
 4 moving data into a computer system's memory, and *writing*, or moving data from
 5 a computer system's memory to another location.

6 Landgraff Decl., p. 0565. An "I/O request" is a "request by an application to read or write a
 7 specified amount of data." EMC's Response, Ex. A (Dictionary of Storage Networking
 8 Terminology). Just as in the claim language, nothing in the dictionary definitions limits which
 9 components of the computer system may contain such applications or issue such requests. In
 10 contrast, consistent with the definitions, the claim language indicates that I/O requests serve to
 11 organize data; that is, data are moved among the system's components regardless of the origin of
 12 the I/O request.

13 Accordingly, I/O requests are "commands, such as reads and writes, used to transfer data
 14 among various components or portions of the components of the computer system."

15 **E. Disputed terms of claim 1 of the '453 patent.**

16 **1. "fully folded array"**

17 The text of claim 1 is set forth below with the key disputed term highlighted in bold type.
 18
 19 In a storage system having $n+1$ disks arranged in a RAID array, a plurality of data
 20 blocks arranged into a plurality of data chunks, a plurality of parity blocks
 21 arranged into a plurality of parity chunks, each parity block associated with n data
 22 blocks in n data chunks, said data chunks and said parity chunks distributed over
 23 said $n+1$ disks, one of said parity chunks and all of said data chunks that are
 24 associated with said parity chunk forming a strip, a method of reorganizing said
 25 data chunks when one of said $n+1$ disks fails, comprising the steps of:

26 detecting the failure of one of said $n+1$ disks;
 27 determining if said failed disk contains all parity chunks;
 28 if said failed disk contains all parity chunks, terminating said method;
 29 if said failed disk contains at least some data chunks, then for each strip
 30 containing a data chunk located on said failed disk, regenerating the data
 31 of said data chunk located on said failed disk and writing said regenerated
 32 data onto said parity chunk associated with said data chunk of said failed
 33 disk to form a **fully folded array**.

34 HP asks the Court to construe this term as: "An organization of a RAID array essentially
 35 that of a RAID-0 organized array except that the disk on which a chunk of data is stored may be
 36 unconventional." EMC proposes:

1 A “fully-folded array” is an array of disks that results after the completion of a
 2 “folding” process. The ‘folding’ process begins after a disk in the array fails. The
 3 data of the failed disk is reconstructed and stored onto the parity chunks of the
 4 active disks in the array. The “folding” is complete when all of the parity chunks
 5 of the active disks have been overwritten with the reconstructed data of the failed
 6 disk.

7 The claimed invention enables continuing access to stored data even when a storage disk
 8 has failed. It does so by reconstructing the failed data and writing them onto portions containing
 9 parity data on other disks. When this occurs, the data are accessible, but the disk storage system
 10 no longer has parity data and thus no longer has redundant data storage. Any data that are lost
 11 thereafter cannot be reconstructed or regenerated. In this state, the data storage device is a “fully
 12 folded array.” The user can continue to access the data until a replacement disk is added, at
 13 which time the data are “unfolded” and restored to their redundant state.

14 While the plain language of the claim does not explain the meaning of “fully folded
 15 array,” the specification provides a sufficient explanation for purposes of claim construction. For
 16 example, the summary of the invention clearly defines a fully-folded array as one in which the
 17 level of organization of data is non-redundant and which contains regenerated data from a failed
 18 disk in the place of previously existing parity data corresponding to that regenerated data. *See*
 19 ‘453:2/35-49.

20 HP’s proposed construction is accurate in the sense that it restates an explanation from
 21 the specification. However, it does not take advantage of the full description provided; the
 22 Court’s construction seeks to define fully the terms proposed by HP. While the fully-folded array
 23 is defined partially by characteristics obtained from the process of folding, EMC improperly
 24 attempts to import the process by which the data are fully folded into the construction. At the
 25 same time, the fully-folded array contains regenerated data in places that previously contained
 26 parity data; the Court will include this property in its construction without importing method or
 27 process limitations.

28 Accordingly, the Court construes this term as “a memory storage array in which the level
 29 of organization of data is non-redundant and which contains regenerated data from a failed disk
 30 in the place of previously existing parity data corresponding to that regenerated data.”

1 F. Disputed term of claims 1 and 7 of the '327 patent.

2 1. **"array state bits"**

3 The text of claim 1 is set forth below with the key disputed term highlighted in bold type.

4 A method of on-line reorganization of data in a storage system having n+1 disks
5 arranged in a RAID array when one of said disks fails while preserving concurrent
6 access to said array by a user application, each of said disks in said storage system
7 having a plurality of data blocks, a plurality of parity blocks, each parity block
8 associated with n data blocks, a plurality of groups of parity block state bits, each
group of parity block state bits indicating that said associated parity block contains
parity information or data or the contents are undefined, and a group of **array**
state bits indicating the condition of the array as "normal," "folding," "fully
folded" or "unfolding," said method comprising the steps of:

- 9 (a) receiving a read request from said user application to read a targeted data block;
- 10 (b) interrogating said **array state bits** to determine the condition of said array;
- 11 (c) if said **array state bits** indicate said array is in said "normal" state, then
proceeding with a read operation;
- 12 (d) if said **array state bits** indicate said array is in said "folding" state, then
proceeding with a "folding" read operation;
- 13 (e) if said **array state bits** indicate said array is in said "fully folded" state,
then proceeding with a "fully folded" read operation; and
- 14 (f) if said **array state bits** indicate said array is in said "folding" state, then
also proceeding with a deliberate process to complete said reorganization
concurrently with user application access to said array.

17 The text of claim 7 is set forth below, with the key disputed term highlighted in bold type.

18 A method of on-line reorganization of data in a storage system having n+1 disks
19 arranged in a RAID array when one of said disks fails while preserving concurrent
20 access to said array by a user application, each of said disks in said storage system
21 having a plurality of data blocks, a plurality of parity blocks, each parity block
22 associated with n data blocks, a plurality of groups of parity block state bits, each
group of parity block state bits indicating that said associated parity block contains
23 parity information or data or the contents are undefined, and a group of **array**
state bits indicating the condition of the array as "normal," "folding," "fully
folded" or "unfolding," said method comprising the steps of:

- 24 (a) receiving a write request from said user application to write to a targeted
data block;
- 25 (b) interrogating said **array state bits** to determine the condition of said array;
- 26 (c) if said **array state bits** indicate said array is in said "normal" state, then
proceeding with a write operation;

- 1 (d) if said **array state bits** indicate said array is in said “folding” state, then
proceeding with a “folding” write operation;
- 2 (e) if said **array state bits** indicate said array is in said “fully folded” state,
then proceeding with a “fully folded” write operation; and
- 3 (f) if said **array state bits** indicate said array is in said “folding” state, then
also proceeding with a deliberate process to complete said reorganization
concurrently with user application access to said array.

6 HP asks the Court to construe this term as: “Bits indicating the state of the array as a
7 whole.” EMC proposes: “Bits stored on each disk in a RAID array that indicate the condition of
8 the RAID array as being ‘normal,’ ‘folding,’ ‘fully-folded,’ and ‘unfolding.’”

9 The claim language reveals that array state bits may indicate the condition of the array as
10 normal, folding, fully folded, or unfolding. Nevertheless, a list of things that an object may
11 indicate does not necessarily define the object. Indeed, the specification acknowledges that array
12 state bits identify other information as well. For example, “[i]n the preferred embodiment, two
13 array state bits identify the missing or replacement member disk and two additional array state
14 bits identify the state of the array.” ‘327:11/67-12/2. Similarly, “[t]he state information for the
15 array as a whole includes the following: (a) Identity of the missing or replacement member disk
16 and (b) Array state. The array state includes (a) ‘normal,’ . . . (b) ‘folding,’ . . . (c) ‘fully folded,’
17 . . . and, (d) ‘unfolding.’” ‘327:11/50-57. Accordingly, looking to the specification to help
18 define the term without importing limitations, the Court concludes that “array state bits” contain
19 information relating to the identity of a missing or replacement disk and the state of the array.

20 EMC’s proposal further requires the array state bits to be “stored on each disk in a RAID
21 array.” However, nothing in the claim or the specification indicates that a location limitation is
22 necessary for proper claim construction.

23 Accordingly, the Court construes this term as “bits that identify either the missing or
24 replacement member disk or the state of the array.”

25 **G. Disputed terms of claims 1 and 3 of the ‘979 patent.**

26 The text of claim 1 is set forth below with the key disputed term highlighted in bold type.

27 A storage system capable of selectively presenting logical units to two or more
28 host computing systems, the storage system comprising:

1 one or more persistent storage devices arranged as logical units;
 2 an array controller controlling and coordinating the operations of said
 persistent storage devices;
 3 a memory accessible by said array controller; and
 4 a **configuration table** stored in said memory, said **configuration table**
 containing one or more entries governing one or more interactions
 between the logical units and the two or more host computing systems,
 wherein said **configuration table** contains an access field entry for each
 logical unit, said access field entry indicating if a particular logical unit
 can be permitted by the array controller to communicate with a particular
 one of said host computing systems.

8 The text of claim 3 is set forth below with the key disputed terms highlighted in bold
 9 type.

10 The system of claim 1, wherein said **configuration table** contains a **host mode**
 11 entry for each logical unit, said **host mode** entry indicating if a particular logical
 unit should communicate with a particular one of said host computing systems
 12 using a predetermined **host mode**.

13 1. **“configuration table”**

14 HP asks the Court to construe this term as: “A data structure used to organize
 15 information regarding the assignment of certain operating parameters of the storage system.”
 16 EMC proposes: “A table organized on a logical unit by-logical unit basis containing, for each
 17 logical unit, one or more data fields related to the logical unit.”

18 The meaning of “configuration table” is unclear from the language of claim 1, which
 19 indicates that a “configuration table” may contain “entries governing one or more interactions
 between the logical units and the two or more host computing systems” and “an access field entry
 20 for each logical unit.” Claim 3 includes a further potential characteristic of a configuration table:
 21 “a host mode entry for each logical unit.”

23 The written description states that “the configuration tables store information relating to
 24 which host connection a particular logical unit should communicate with, any LUN numbers
 offsets employed by the logical unit to communicate with a particular host, and any special ‘host
 25 modes’ employed by the array controller in communicating with a particular host.” ‘979:3/64-
 26 4/2. Thus, the claim language and the specification indicate that the configuration table contains
 27 data that govern the interactions between logical units and host computing systems. The

specification and the claims include examples of the specific types of data contained in a configuration table, including “access fields,” “offset fields,” and “host mode fields.” *See e.g.*, ‘979:4/3-36 & 5/21. However, such examples need not be imported as limitations for the purposes of defining this particular term. Indeed, claim 1 explicitly does not include an “offset field entry” limitation or a “host mode entry” limitation (these are found in dependent claims 2 and 3). Using the doctrine of claim differentiation as a guide to claim construction, *Laitram Corp. v. Rexnord, Inc.*, 939 F.2d 1533, 1538 (Fed. Cir. 1991), it would be inappropriate for the Court to specify particular types of entries in its construction of “configuration table.”

HP’s proposal also is ambiguous and potentially is too broad. Under HP’s construction, the configuration table would contribute to organization of “information regarding the assignment of certain operating parameters of the storage system.” However, the claim and the specification indicate that the configuration table is limited to data that govern the interactions between the logical units and the host computing systems. It is unclear from the intrinsic evidence what “assignment of certain operating parameters” means. According to a dictionary cited by HP itself, “a table is a data structure used to organize information.” Langraff Decl., Separate Appendix, vol. 7, p. 1221. Notably, HP concedes that “the configuration table must specify information that governs aspects of the interaction between the logical units of the storage system and the host computing systems connected to the storage system.” HP’s Opening Brief, pp. 39-40.

At the same time, EMC’s proposal is too narrow. While certain data relating to logical units obviously are contained within the configuration table, nothing in the claim language or the specification requires limiting the configuration table to a certain internal organization (“a table organized on a logical unit by-logical unit basis”).

Accordingly, the Court construes this term as “a structure organizing data that governs the interactions between logical units and host computing systems.”

2. **“host mode”**

HP asks the Court to construe this term as: “An operational state that facilitates communication with a particular host.” EMC proposes: “A mode, unique to a particular

1 operating system, that specifies the manner in which an array controller should communicate
 2 with a host computer that is running said operating system.”

3 Although the meaning of “host mode” is unclear from the claim language alone, a “host
 4 mode entry,” according to the Court’s construction of the preceding term, must be “data that
 5 governs the interactions between logical units and host computing systems.” Indeed, according
 6 to the claim, a “host mode entry” may indicate “if a particular logical unit should communicate
 7 with a particular host computing system” and the host mode is predetermined.

8 For clarification, the Court must turn to the specification, which describes embodiments
 9 in which the host mode relates to the host’s operating system. For example, “unique host
 10 ‘modes’ [are] adapted to be employed uniquely to each host operating system to which the
 11 logical unit . . . is coupled. These unique host modes permit the array controller to be adaptable
 12 to handle different idiosyncrasies of particular host operating systems.” ‘979:4/31-36. Although
 13 in one embodiment the host mode field “corresponds to the name of the operating system running
 14 on the particular host,” ‘979:7/37-38, the host mode entry need not be limited to data concerning
 15 only a particular host or operating system. “For example, if a particular host is found to operate
 16 more efficiently or more smoothly in a particular mode, then based on the host operating system
 17 type, [such information can be] stored in the configuration table.” ‘979:6/45-51. That is,
 18 although host mode entries contain information related to idiosyncrasies of a particular host,
 19 nothing in the intrinsic evidence requires each host to have a distinct or unique idiosyncrasy or
 20 associated host mode entry. Groups of hosts could share the same idiosyncrasy or host mode
 21 entry. Thus, a host mode entry could contain information about a particular operating system or
 22 the “operating system type.”

23 The specification does not limit host mode entries solely to information relating to
 24 operating systems or operating system types. For example, the configuration table may include a
 25 data field such as “any special host modes for the particular host adapter.” ‘979:5/11-12. The
 26 “host’s adapter connects to the array controller.” ‘979:5/41-42. This adapter contains
 27 identification information of the host, which is stored in the configuration table. ‘979:5/42-45.
 28 In this way, access to of the host to the logical units can be controlled. Thus, “host mode” data

1 are additional information about the host, such as any of its “idiosyncrasies,” ‘979:6/42-51—and
 2 not only the host’s operating system.

3 Accordingly, a “host mode entry” may be an entry of data documenting the particular or
 4 idiosyncratic characteristics or the operational state of a host or group of hosts. This is consistent
 5 with a dictionary definition of “mode” as an “operational state of a computer or a program.”
 6 Landgraff Decl., Separate Appendix, vol. 7, p. 1219 (Microsoft Computer Dictionary). Such
 7 information may be used in accordance with the invention to facilitate communication with a
 8 particular host (part of HP’s proposal) or to specify the manner in which an array controller
 9 should communicate with a host computer (part of EMC’s proposal). It follows that the term
 10 “host mode” relates only to the state of the host, not the purpose or use of the “host mode” data.
 11 Indeed, HP concedes that a “mode” is an “operational state of a computer.” HP’s Opening Brief,
 12 p. 40.

13 Accordingly, the Court construes this term as “the particular or idiosyncratic operational
 14 state of a host or group of hosts.”

15 **H. Disputed terms of the ‘801 patent.**

16 The text of claim 1⁵ is set forth below with the key disputed terms highlighted in bold
 17 type.

18 In a computer system including a processor and a peripheral device, the processor
 19 executing an application program which interacts With [sic] the peripheral device
 20 by sending peripheral device requests to operating system routines which send
 21 corresponding peripheral device commands to the peripheral device in response to
 22 the peripheral device requests, wherein the peripheral device is capable of
 23 replacing resident **microcode** with new **microcode** by download by the
 24 application program, and comprises:

25 a non-volatile memory containing the resident **microcode**;

26 circuitry for receiving peripheral device commands sent by the operating
 27 system in response to peripheral device requests from the application
 28 program; and

29 a resident processor, coupled to the non-volatile memory and receiving
 30 circuitry, and responsive to the resident **microcode**, comprising:

5 Although the Court refers to representative claims 1 and 20, it will construe the terms as
 used throughout the patent.

1 a detector of an **initiator peripheral device command** from the operating
 2 system corresponding to an initiator peripheral device request from the
 3 application program;

4 a detector of a **transfer peripheral device command** from the operating
 5 system corresponding to a transfer peripheral device request from the
 6 application program including the new **microcode**, received while in a
 7 **waiting state**; and

8 circuitry, responsive to the initiator command detector and the transfer
 9 command detector, for entering the **waiting state** when an **initiator**
 10 **peripheral device command** is detected, and for transferring the new
 11 **microcode** from the receiving circuitry into the non-volatile memory
 12 and restarting the operation of the peripheral device when a **transfer**
 13 **peripheral device command** is detected.

14 The text of claim 20 is set forth below with the key disputed terms highlighted in bold
 15 type.

16 A method of operating a peripheral device responsive to peripheral device
 17 commands corresponding to peripheral device requests from an application
 18 program, to replace resident **microcode** with new **microcode** by download
 19 comprising the steps of:

20 responding to an **initiator peripheral device command**, corresponding to an
 21 initiator peripheral device request by the application program, by entering
 22 a **waiting state** ready to receive new **microcode**; and

23 if in the **waiting state**, responding to a **transfer peripheral device command**,
 24 including the new **microcode**, corresponding to a transfer peripheral
 25 device request by the application program, by replacing the resident
 26 **microcode** with the new **microcode** and restarting the operation of the
 27 peripheral device.

28 **1. “microcode”**

29 EMC asks the Court to construe this term as: “One or more programs for a peripheral
 30 device that can be stored in the non-volatile memory of the peripheral device.” HP initially
 31 proposed: “Program instructions embedded in the internal circuitry of a microprocessor.” Thus,
 32 the parties initially disputed where the microcode is stored. However, HP changed its proposal in
 33 its opposition brief to: “Program instructions that directly control a microprocessor, stored in
 34 internal circuitry that is non-volatile memory directly connected to that microprocessor.”⁶ HP

28 ⁶ EMC has filed a “Motion to Preclude HP from Asserting New Claim Construction
 29 Positions in its Opposition Brief.” The Court has considered the motion on a term-by-term basis.

1 thus concedes that the microcode need not be stored in the microprocessor. Plaintiffs' and
 2 Counterclaim Defendants' Opposition Claim Construction Brief on Key Terms, p. 11.

3 The claims do not define microcode; they merely indicate that microcode resides in non-
 4 volatile memory and may be replaced. The specification, however, provides useful insight.

5 Peripheral devices have "special processors installed in them, called resident processors."

6 '801:1/14-15. "These resident processors execute programs, called microcode in the remainder
 7 of this application, which are generally stored in a non-volatile memory in the peripheral device."

8 '801:1/16-19. Thus, in its most general sense, "microcode" is a program executed by a
 9 peripheral device's resident processor. The remainder of the specification is consistent with this
 10 interpretation. *See, e.g.*, '801:2/33-35 ("resident processor . . . is responsive to the resident
 11 microcode"), 2/50-57, 5/47-51, 6/35-37, & 7/46-48.

12 Although the disputed claims do not include the limitation of "control," HP proposes that
 13 microcode "controls" the peripheral device. HP is correct in implying that "microcode" does not
 14 include all possible types of computer code or programs. Indeed, the patent itself is entitled
 15 "Method and Apparatus for Replacing Resident Peripheral Device Control Microcode . . .," and
 16 the written description nearly always modifies "microcode" with the word "control." EMC does
 17 not necessarily dispute that microcode is a limited subset of all possible computer programs or
 18 code: it concedes that microcode is a program "for a peripheral device," that is, code used to
 19 assist a peripheral device in functioning. The specification uses the word "microcode" in a
 20 manner consistent with this interpretation. *See, e.g.*, '801:1/25-29 ("In order to upgrade
 21 performance of, or provide new features for, such peripheral devices, the control microcode may
 22 be updated via a download of new control microcode."), 5/47-50 ("The resident processor 106

23 Whether the motion should be granted or additional briefing from EMC is appropriate in light of
 24 the belatedness of HP's proposals depends upon whether EMC has had fair opportunity to
 25 present its position. The Court is satisfied that EMC had such an opportunity given the length
 26 and thoroughness of the hearing. Moreover, because claim construction is not only a matter of
 27 resolving disputes between litigants but also has a potentially preclusive effect with respect to
 28 litigation involving litigants not party to the present suit, the Court will consider HP's new
 positions to the extent that they assist the Court in construing the claim terms properly.
 Accordingly, the motion will be denied. For the sake of brevity, the Court generally will not
 restate HP's initial proposals. For the initial proposals see EMC's reply brief.

1 operates in a known manner, according to a control program stored in the control microcode in
 2 the non-volatile memory 108, to execute the disk control command.”). Similarly, a dictionary
 3 definition provided by both parties defines microcode in part as the “lowest-level instructions
 4 that directly control a microprocessor.” Declaration of Mark R. Weinstein in Support of EMC’s
 5 Opening Claim Construction Brief, Ex. 3 (Random House Webster’s Computer & Internet
 6 Dictionary). While the Court will not import the word “control” into the claim, it notes that a
 7 person of ordinary skill in the art would understand “microcode” to mean a limited set of
 8 programs that are used specifically *for* a peripheral device.

9 With respect to the parties’ contentions regarding the location of the microcode, it is true
 10 that microcode generally is stored in non-volatile memory in a peripheral device. However,
 11 “microcode” need not be defined according to its location, and nothing in the claim indicates that
 12 such a limitation must be added. Moreover, the importation of such a limitation would be
 13 inconsistent with the written description. For example, microcode sometimes is transferred from
 14 a central computer system for replacement purposes. *See* ‘801:1/26-29.

15 Similarly, nothing in the claim language or specification requires the microcode to control
 16 a microprocessor “directly” or that the non-volatile memory be “directly” connected to a resident
 17 microprocessor. Although microcode sometimes may “control” a resident processor, *see*
 18 ‘801:4/21-23, it is unclear what is meant by “direct” connection or control. The purpose of claim
 19 construction is to clarify the claim language, not to add ambiguous terms. Moreover, the
 20 invention appears to teach away from “direct connection” of the resident microprocessor and the
 21 non-volatile memory. *See* ‘801:Figure 1. Accordingly, the Court construes this term as “one or
 22 more programs for a peripheral device that can be stored in the non-volatile memory of the
 23 peripheral device and executed by a peripheral device’s resident processor.”

24 **2. “waiting state”**

25 EMC asks the Court to construe this term as: “State of operation of the peripheral device
 26 wherein the peripheral device is ready to receive new microcode.” HP proposes (in its opposition
 27 brief): “A state in which a peripheral device stands ready to receive new microcode.” The
 28 parties concede that their constructions are the same. Transcript of April 5, 2004 Hearing, 82:3-

1 4; 87:12-16.

2 The meaning of “waiting state” is unclear from the claim language. The claim indicates
 3 that a detector of transfer peripheral device command receives the command and new microcode
 4 while in a waiting state and that there is circuitry for entering the waiting state in response to the
 5 initiator and transfer command detectors. The specification states:

6 When an initiator command is detected, a waiting state is entered in which the
 7 peripheral device is held ready to receive new microcode. The resident processor
 8 also includes a detector for a transfer peripheral device command, which includes
 9 the new microcode, received while the peripheral device is in the waiting state.

10 ‘801:2/36-42; *see also* ‘801:2/53 & 6/39-44. Thus, the “waiting state” is equivalent to the state
 11 of the peripheral device when it is ready to receive new microcode. Accordingly, the Court
 12 construes this term as “a state of the peripheral device in which it is ready to receive new
 13 microcode.”

14 **3. “initiator peripheral device command”**

15 EMC asks the Court to construe this term as: “A command that signals to the peripheral
 16 device the beginning of the process of replacing resident microcode.” EMC also is willing to add
 17 the following to its proposal: “This command is sent by the operating system to a peripheral
 18 device in response to an initiator peripheral device request sent by an application program to the
 19 operating system, and causes the peripheral device to enter a waiting state ready to receive new
 20 microcode.” HP proposes (in its opposition brief):

21 A command sent by the operating system to a peripheral device in response to an
 22 initiator peripheral device request sent by an application program to the operating
 23 system. The command causes the peripheral device to enter a waiting state ready
 24 to receive new microcode. The command is separate and functionally different
 25 from a subsequently received transfer peripheral device command.

26 The initiator peripheral device command is one that an operating system sends to
 27 the peripheral device in response to a routine legitimately called by an application
 28 program (e.g., a write or read request), as opposed to a special command for the
 29 peripheral device to which an application program would not normally have
 30 access.

31 Although the claim language does not clearly define the term, it indicates that the
 32 peripheral device enters the waiting period in response to the initiator peripheral device
 33 command. Additionally, the initiator peripheral device command corresponds to an initiator
 34

1 peripheral device request sent from the application program—executed by the computer system
 2 processor—to the operating system routine. The written description further explains:

3 The operating system provides operating system routines, legitimately accessible
 4 to an application program, for providing such access. In response to a call by an
 5 application program to such an operating system routine, the operating system will
 6 send a specific peripheral device control command . . . to the peripheral device to
 7 perform the requested access.

801:2/4-10.

9 It is not clear from the claim language whether the transfer and initiator commands must
 10 be “separate and functionally different.” The claim language indicates that the peripheral device
 11 enters a waiting state in response to the initiator peripheral device command and that transfer of
 12 microcode to the peripheral device cannot occur in the absence of a transfer peripheral device
 13 command, strongly suggesting that there are two distinct commands. The claim does not state
 14 that microcode may be transferred in response to the initiator peripheral device command or that
 15 the transfer peripheral device command can induce a waiting state. Thus, it appears from the
 16 claim language alone that the two commands are distinct. The specification strongly suggests
 17 that the transfer and initiator commands must be different commands. For example, the resident
 18 processor of the peripheral device

19 detect[s] a first disk control command, called the intiator command . . . , in
 20 response to which the disk drive is placed in a state waiting to receive new control
 21 microcode. When in this waiting state, a second disk control command, called the
 22 transfer command . . . , is detected which contains the new control microcode for
 23 the disk drive.

24 ‘801:6/37-44. Similarly, “[w]hen the initiator command is detected, no transfer of data is
 25 performed to the disk drive platters 114. Instead, the disk drive enters a state ready to receive
 26 new microcode.” ‘801:8/63-65. Finally, the prosecution history indicates that the patentee
 27 distinguished the invention from prior art that disclosed a “system which enters the waiting state
 28 when a single . . . command containing the complete replacement microcode is being received
 and processed,” and the patentee referred to the present invention as using a “two command
 sequence . . . to download replacement microcode.” Declaration of Steven E. Derringer in
 Support of Plaintiffss and Counterclaim Defendants’ Opposition Claim Construction Brief on
 Key Terms (“Derringer Decl.”), Separate Appendix, vol. 1, pp. 0170, 0174-75. The patentee’s

1 description is consistent with the claim language and the specification. Accordingly, in the
 2 claimed invention, the peripheral device must receive two distinct commands in order to receive
 3 microcode.

4 HP also proposes language excluding a “special command.” The basis for the limitation
 5 is found in the prosecution history, in which the patentee distinguished the claimed invention
 6 from the prior art by stating that it allowed for updating of resident microcode via operating
 7 system routine, in contrast to special commands that would require shutting down the computer
 8 system. This statement is consistent with a construction that includes reference to operating
 9 system routines that respond to application programs, in that a special command, by definition,
 10 does not utilize this process. *See* Derringer Decl., Separate Appendix, Vol.1, p. App. 0170 (The
 11 “operating system forbids access to [special] commands by the application program.”). Thus, the
 12 inclusion of reference to “special commands” is unnecessary.

13 Accordingly, the Court construes this term as “a command sent by the operating system to
 14 a peripheral device in response to an initiator peripheral device request sent by an application
 15 program to the operating system, and which causes the peripheral device to enter a waiting state
 16 ready to receive new microcode.”

17 **4. “transfer peripheral device command”**

18 EMC asks the Court to construe this term as: “A command that transfers new microcode
 19 to the peripheral device.”⁶ HP proposes (in its opposition brief):

20 A command sent by the operating system to a peripheral device in response to a
 21 transfer peripheral device request sent by an application program to the operating
 22 system. The command is received by the peripheral device while the device is in
 23 a waiting state, and it causes the peripheral device to download new microcode.
 24 The command is separate and functionally different from a previously received
 25 initiator peripheral device command.

26 ⁶ EMC stated at the Claim Construction Hearing on April 5, 2004, that it is willing to add
 27 language similar to the language added for the prior term. Transcript of April 5, 2004 Hearing,
 28 88:21-24. Obviously, EMC did not intend to adopt the exact language, and thus the Court must
 alter it so that it is consistent with the meaning of this term. The Court assumes, only for the
 purposes of assisting in the discussion, that EMC intended to add the following to its proposal:
 “This command is sent by the operating system to a peripheral device in response to a transfer
 peripheral device request sent by an application program to the operating system, and causes the
 peripheral device to receive new microcode.”

1 The transfer peripheral device command is one that an operating system sends to
 2 the peripheral device in response to a routine legitimately called by an application
 3 program (e.g., a write or read request), as opposed to a special command for the
 4 peripheral device to which an application program would not normally have
 5 access.

6 Although the claim language does not clearly define the term, it indicates, consistent with
 7 the specification, that the peripheral device receives the transfer peripheral device command
 8 while it is in the waiting state. The new microcode is transferred to the peripheral device when it
 9 detects the transfer peripheral device command. Additionally, the transfer peripheral device
 10 command corresponds to a transfer peripheral device request sent from the application
 11 program—executed by the computer system processor—to the operating system routine. The
 12 written description further explains:

13 The operating system provides operating system routines, legitimately accessible
 14 to an application program, for providing such access. In response to a call by an
 15 application program to such an operating system routine, the operating system will
 16 send a specific peripheral device control command . . . to the peripheral device to
 17 perform the requested access.

18 ‘801:2/4-10.

19 HP’s proposal includes reference to the initiator peripheral device command. As
 20 discussed above, the prosecution history indicates that the patentee intended to claim an
 21 invention that uses a “two command sequence . . . to download replacement microcode.”
 22 Deringer Decl., Separate Appendix, vol. 1, pp. 0170, 0174-75. This interpretation is consistent
 23 with the claim language and the specification, both of which describe a sequence of actions
 24 through which the peripheral device may receive microcode. That is, the peripheral device
 25 cannot receive microcode as part of the transfer peripheral device command prior to entering the
 26 waiting state in response to an initiator peripheral device command. HP also proposes language
 27 excluding a “special command.” As stated above, a construction that includes access by an
 28 application program necessarily precludes the possibility of special commands.

29 Accordingly, the Court construes this term as “a command received by a peripheral
 30 device—after it has entered the waiting state in response to the initiator peripheral device
 31 command—from an operating system in response to a transfer peripheral device request sent by
 32 an application program to the operating system, and which causes the peripheral device to receive

1 new microcode.”

2 **I. Disputed terms of claim 1 of the ‘497 patent.**

3 The text of claim 1 is set forth below with the key disputed terms highlighted in bold
4 type.

5 A method for identifying members of a logical set of mass storage devices,
6 comprising the steps of:

7 defining a **global identifier** for said logical set which serves as a reference for
8 identifying mass storage devices included in said logical set;

9 assigning **membership signatures** to a plurality of mass storage devices,
10 wherein said **membership signatures** comprise information substantially
11 matching said **global identifier** which identifies said plurality of mass
12 storage devices as members of said logical set; and

13 when a member of said logical set undergoes a **change in membership**
14 **status**, defining a new **global identifier** for said logical set different from
15 said **global identifier** which thereafter serves as a reference for identifying
16 members of said logical set, and assigning new **membership signatures**
17 to other ones of said plurality of mass storage devices that remain
18 members of said logical set, wherein said new **membership signatures**
19 comprise different information substantially matching said new **global**
20 **identifier** which identifies said other ones of said plurality of mass storage
21 devices as members of said logical set.

22 **1. “global identifier”**

23 EMC asks the Court to construe this term as: “A ‘global identifier’ identifies the mass
24 storage devices in a logical set.” HP proposes (in its opposition brief and at the claim
25 construction hearing, *see* Transcript of April 6, 2004 Hearing, 151:14-22): “A value or data
26 structure stored in the array controller that is compared to membership signatures in order to
27 identify which mass storage devices are members of a logical set of mass storage devices.”
28 Essentially, then, the parties dispute whether the term has a location limitation.

29 The meaning of “global identifier” can be ascertained from the claim language, which
30 indicates that a “global identifier” serves as a reference for identifying mass storage devices
31 included in a logical set. HP incorrectly argues that this is not a definition of “global identifier”
32 but instead describes a *use* of a “global identifier.” The construction in fact is a definition: the
33 global identifier is a reference with which each storage device in a logic set identifies. In other
34 words, knowledge of the global identifier allows the system to identify members of the set.

1 Nothing in the intrinsic evidence indicates that this term must be defined in part
 2 according to its location, nor is the term defined according to a process or method of comparing
 3 it to something else or according to its purpose. Accordingly, HP's proposed limitation requiring
 4 that the global identifier be stored in the array controller is inappropriate.

5 The Court construes this term as "a reference for identifying mass storage devices
 6 included in a logical set."

7 **2. "membership signatures"**

8 EMC asks the Court to construe this term as: "Information which identifies a mass
 9 storage device as a member of a logical set and substantially matching the global identifier." HP
 10 proposes (in its opposition brief and at the claim construction hearing, *see* Transcript of April 6,
 11 2004 Hearing, 161:3-10): "A value or data structure stored either in the array controller or on the
 12 mass storage device, reflecting the device's status as an active member of a logical set of mass
 13 storage devices by substantially matching the global identifier."

14 The claim language indicates that membership signatures comprise information
 15 substantially matching the global identifier. The specification indicates that a membership
 16 signature is assigned "to each of the physical mass storage devices in the set [and it] identifies the
 17 physical mass storage device as a member of the set." '497:2/41-44. The set may be a logical set
 18 comprised of portions of the physical disk drives. *See* '497:3/12-16. Additionally, "[t]he
 19 signatures for all members are identical, or at least substantially identical (e.g., the signature may
 20 include a member number that differs from drive to drive.)" '497:3/16-19.

21 HP's proposed additional limitations are not supported by the intrinsic evidence. The
 22 patentee expressly claimed a location limitation in dependent claim 8, strongly suggesting in light
 23 of the doctrine of claim differentiation that there is no location limitation in the independent
 24 claim. Moreover, nothing in the claim supports the importation of a limitation requiring the
 25 comparison of the global identifier to the membership signature in order to determine whether its
 26 status is "active." Instead, the members of a set are assigned membership signatures and the set
 27 is given a substantially matching global identifier.

28 Accordingly, the Court construes "membership signatures" as "information that identifies

1 physical mass storage devices or portions of physical mass storage devices as members of a set
 2 and which substantially matches the global identifier of the set.”

3 **3. “change in membership status”**

4 EMC asks the Court to construe this term as either: “change in membership status” or “a
 5 change in status resulting from an act affecting the member, such as the member being removed
 6 from the set.” HP proposes (in its opposition brief): “Exclusion of a physical mass storage
 7 device from a logical set of mass storage devices (and, in certain circumstances, the writing of
 8 data onto the remaining physical mass storage devices in the logical set) and updating the
 9 membership of the remaining physical mass storage devices in the logical set.”

10 The claim language indicates that when there is a change in membership status, “other
 11 ones of said plurality of mass storage devices that remain members of said logical set” are
 12 assigned new membership signatures and the set is assigned a new substantially matching global
 13 identifier. The specification states that “[w]hen a member of the set undergoes a change in
 14 membership status the [sic] (e.g., removal of a physical mass storage device), a new membership
 15 signature is assigned to each remaining physical mass storage device in the set.” ‘497:2/45-49.
 16 Thus, removal is only one possibility of status change. Indeed, the written description states:
 17 “whenever *any event* occurs to change the membership of any of the drives in the set . . .”
 18 ‘497:3/20-22 (emphasis added). Removal or addition of a storage device occurs in the preferred
 19 embodiment. ‘497:2/23-26. Finally, dependent claim 2 limits the change in membership status
 20 to “the removal of said member from said logical set.” Importing that limitation into
 21 independent claim 1 would render claim 2 superfluous. The Court therefore will not import
 22 “exclusion of a storage device” into the construction of this term. Similarly, there is nothing in
 23 the claim that requires writing data onto or updating the membership of the remaining members
 24 of the storage memory devices in the set (a limitation that appears to be the subject of claim 5).
 25 Although such acts may result from a change in membership, they are not necessary to a
 26 definition of the term.

27 Accordingly, the Court construes “change in membership” as “a change in status resulting
 28 from an act affecting a member of the set, such as a member being added or removed from the

1 set."

2 **J. Disputed terms of claims 1 and 13 of the '184 patent.**

3 The text of claim 1 is set forth below with the key disputed term highlighted in bold type.

4 In a digital data processing system including at least one **local digital data**
 5 **processing system**, each said at least one local system including processor means
 6 for processing data items in response to instructions, a universal addressing
 7 system for addressing said data items comprising:

8 (1) universal logical memory means accessible to each said at least one local
 9 system for storing and providing said data items,

10 (2) memory organization means for organizing said universal memory means into
 11 objects uniquely and permanently identified by unique identifiers, each said data
 12 item being associated with a said object and addressable by means of a logical
 13 address specifying the unique identifier for said object with which said data item
 14 is associated and an offset specifying the locations of said data item in said
 15 associated object

16 (3) means in said processor means for providing memory operation specifiers in
 17 response to said instructions, said memory operation specifiers each including a
 18 logical address and a memory command specifying a memory operation, and

19 (4) memory operation means responsive to a memory operation specifier for
 20 accessing the data item and performing the memory operation specified by the
 21 memory command in said memory operation specifier on the data item specified
 22 by the logical address in said memory operation specifier.

23 The text of claim 13 is set forth below with the key disputed terms highlighted in bold
 24 type.

25 A digital data processing system comprising:

26 (1) universal logical memory means for storing and providing items of data
 27 including instructions,

28 (2) **memory organization means for organizing said universal logical memory**
means into objects uniquely and permanently identified by unique
identifiers, each said data item being associated with a said object and
 29 addressable by means of a logical address specifying the unique identifier of the
 30 object with which said data item is associated and an offset specifying the location
 31 of said data item in said associated object, and [sic]

32 (3) at least one **local digital data processing system** having access to said
 33 universal logical memory means and including processor means for processing
 34 said data items and providing memory operation specifiers in response to said
 35 instructions, said memory operation specifiers each including a logical address
 36 and a memory command specifying a memory operation, and

37 (4) memory operation means responsive to a memory operation specifier for
 38 accessing the data item and performing the memory operation specified by the

1 memory command on the data item specified by said logical address.

2 **1. “local digital data processing system”**

3 EMC asks the Court to construe this term as: “A digital processing system (*i.e.*, a
4 computer system) restricted to a particular area.” HP proposes (in its opposition brief): “A
5 computer system comprising a processor that is directly connected to a main memory and an
6 input/output system.”

7 The claim language indicates that the local digital data processing system is part of a
8 digital data processing system and that it includes a processor for processing data. Although the
9 term does not appear in the written description, the specification does discuss a “local system,”
10 *see, e.g.*, ‘184:abstract, which appears to be synonymous with the term. “The digital computer
11 system in which the present invention is employed includes at least one local system including a
12 processor and a universal memory accessible to the local system for storing data including
13 instructions.” ‘184:2/67-3/3. Thus, as described in the specification, the digital computer system
14 is comprised of a universal memory and local digital computer systems that can access that
15 universal memory. The claims acknowledge this view with additional limitations (“universal
16 logical memory means accessible to each said at least one local system” and “local digital data
17 processing system having access to said universal logical memory means”). The claims also
18 include a limitation requiring the local digital processing system to include a “processor means.”
19 These limitations would be redundant if they formed part of the construction of the term.

20 The claim language requires “at least one” local system. However, it is instructive to
21 refer to the ‘5602 patent, which is incorporated by reference, to understand the purpose of
22 including “local systems.” That patent describes an advantage of the invention as

23 providing a flexible internal system structure capable of performing multiple,
24 concurrent operations, comprising a plurality of separate, independent processors,
25 each having a separate microinstruction control and at least one separation and
independent port of a central, independent communications and memory
processor comprised of a plurality of independent processors capable of
performing multiple, concurrent memory and communications operations.

26 ‘602:4/3-11. Thus, the purpose of having “local” systems is that there could be multiple “local”
27 systems. While claims 1 and 13 specifically include the possibility of having a single such
28

1 system, it is the nature of a “local system” that there is the possibility of more than one distinct
 2 system. Accordingly, the word “local” signifies that the invention may include multiple,
 3 discretely located, and distinct systems.

4 Finally, nothing in the specification or claims of this patent or the ‘5602 patent requires a
 5 processor to be “directly connected to a main memory and an input/output system.” These
 6 proposed limitations are part of a described embodiment. Moreover, it is unclear what “directly”
 7 connected to would mean in the context of the claim.

8 Accordingly, the Court construes this term as “a digital data processing system restricted
 9 to a particular area and physically and logically distinct from any other local digital data
 10 processing system if one exists.”

11 **2. “memory organization means for organizing said universal logical memory
 12 means into objects uniquely and permanently identified by unique
 13 identifiers”**

14 EMC argues that despite the use of the word “means” this is not a means plus function
 15 claim, and it therefore asks the Court to construe this term as: “Memory organization means
 16 organizes the memory, which can include disk drives, into objects.” Alternatively, should the
 17 Court construe the claim as a means plus function claim, EMC contends that the function of
 18 “memory organization means” is “organizing universal memory means into objects uniquely and
 19 permanently identified by unique identifiers” and its corresponding structure is “an object
 20 management system.” HP proposes (in its opposition brief): “A means for organizing (i.e.,
 21 arranging in a desired pattern or structure) universal memory means into objects uniquely and
 22 permanently identified by unique identifiers, having structures that are the same as or
 23 substantially equivalent to the Object Management System described in U.S. Patent No.
 24 4,455,602 at Col. 20:32-21:16, Col. 49:60-Col. 51:5, Col. 62:10-Col. 64:52, Col. 356:64-Col.
 25 364:45, Col. 367:45-Col. 369:7, and all corresponding figures including but not limited to Figure
 26 2, Figure 101, Figure 106, Figure 106A and Figure 107.”

27 The Court agrees with HP that this is a means plus function claim. The claim clearly
 28 states a function: “organizing said universal logical memory means into objects uniquely and

1 permanently identified by unique identifiers.” There is no clearly recited structure, material, or
 2 act in the claim for performing this function. Thus the Court must look to the specification for
 3 disclosed structure, material, or acts that are clearly linked or associated with this function.
 4 Further complicating the matter, the function includes reference to another “means”: the
 5 “universal logical memory means.” This refers to the preceding subparagraph, which states that
 6 the function of the “universal logical memory means” is “storing and providing said data items.”

7 Because the ‘184 specification does not provide corresponding structure, material, or acts,
 8 the Court must look to the ‘5602 patent, which was incorporated by reference. The Federal
 9 Circuit, in *Atmel Corp. v. Info. Storage Devices, Inc.*, 198 F.3d 1374 (Fed. Cir. 1999), held that a
 10 patentee may use sources that were incorporated by reference to define the disclosed structure,
 11 material, or acts of a means plus function claim if a person having ordinary skill in the art would
 12 do so. Nonetheless, “Paragraph 6 does not contemplate the kind of open-ended reference to
 13 extrinsic works that ¶ 1, the enablement provision, does.” *Id.* at 1382. Thus, the ‘5602 patent
 14 disclosure only can be used to help define the structure, material, or acts disclosed in the ‘184
 15 patent if a person having ordinary skill in the art would use that disclosure to help define the
 16 structure, material, or acts. Moreover, to “incorporate material by reference, the host document
 17 must identify with detailed particularity what specific material it incorporates and clearly indicate
 18 where that material is found in the various documents.” *Advanced Display Sys., Inc. v. Kent*
 19 *State Univ.*, 212 F.3d 1272 (Fed. Cir. 2000). The ‘184 patent clearly indicates expressly that the
 20 entire ‘5602 patent is incorporated by reference.

21 The “universal logical memory means” is shown as MEM 112 in Figures 1 and 18 and is
 22 described at Column 36, line 42 through Column 38, line 41. It also is shown, more specifically,
 23 as MEM 10112 in Figure 101. In particular, “MEM 112 is an intelligent, prioritizing memory.”
 24 ‘5602:36/52. “Major elements of MEM 112 are Main Store Bank (MSB) 1810, a Bank
 25 Controller (BC) 1814, a Memory Cache (MC 1816), a Field Interface Unit (FIU) 1820, and
 26 Memory Interface Controller (MIC) 1822.” ‘5602:36/45-49. MSB 1810 “is comprised of one or
 27 more Memory Array Cards (MAs) 1812” of various capacity. ‘5602:37/13-14 & 37/22-23. “MC
 28 1816 includes a bypass write path” and a “cache write-back path.” ‘5602:37/55 & 37/58. Data

1 transfer between internal components and to and from external components occurs through
 2 buses. *See, e.g.*, '5602:38/12-27. MEM 10112 may use "a single, large memory" or "physically
 3 separate memories." '5602:53/62 & 53/65-66. MEM 10112 is "structured into frames." *See,*
 4 *e.g.*, '5602:63/56-57.

5 The parties agree that the corresponding structure of the "memory organization means for
 6 organizing said universal logical memory means into objects uniquely and permanently identified
 7 by unique identifiers" is an "object management system." HP, however, seeks to define this term
 8 further by reference to the specific embodiments disclosed in the '5602 patent. The Court agrees
 9 that "object management system" is insufficient description of structure, material, or act to define
 10 this term. The specification describes the system in more detail in various places. For example,
 11 Figure 2 shows the organization of MEM 112. The figure is described in column 20:
 12 information stored on MEM 112 is organized into "Objects," each of which "is uniquely
 13 identified by a serial number referred to as a Unique Identifier (UID)." '5602:20/52-54. "UIDs
 14 provide an addressing base common to all CS 101 systems which may ever exist, through which
 15 any Object ever created may be permanently and uniquely identified." '5602:20/60-63.
 16 "[I]nformation within MEM 112 is located through MEM 112 Physical Addresses identifying
 17 particular locations within MEM112's memory space." '5602:21/44-47; *see also* '5602:50/38 &
 18 53/17. Physical Addresses contain a "Frame Number (FN) field, a Displacement (D) field, and a
 19 Length (L) field. Together the Frame Number field and the Displacement field specify the
 20 location in MEM 112 containing the data, and the Length field specifies the length of the data."
 21 '5602:32/54-58 & Fig. 12.

22 Although the specification provides a more detailed description, a sufficient definition of
 23 structure, material, or acts necessary to perform the claimed function can be determined from this
 24 description alone. Moreover, this description, while providing sufficient corresponding structure,
 25 is general enough to avoid rendering the independent claims redundant. *See, e.g.*, *Wenger Mfg.,*
 26 *Inc. v. Coating Mach. Sys., Inc.*, 239 F.3d 1225, 1234 (Fed. Cir. 2001).

27 Accordingly, the Court construes this term as follows: The function is: "organizing said
 28 universal logical memory means into objects uniquely and permanently identified by unique

1 identifiers.” The corresponding structure, material, or act is: “The act of storing information into
 2 a universal logical memory means as described on page 48, line 21 through page 49, line 4 of this
 3 Order, in which information is organized by an Object Management System as described on page
 4 49, line 10 through page 49, line 21 of this Order.”

5 **K. Disputed terms of claim 1 of the ‘756 patent.**

6 The text of claim 1 is set forth below with the key disputed terms highlighted in bold
 7 type.

8 A method for automatically **recovering** a computer system following discovery of
 a fault condition, comprising the steps of:

9 providing a computer system having self-testing and self-diagnosing
 10 capability;

11 automatically testing the computer system;

12 automatically identifying the presence of one or more faulted components
 from said step of testing; and

13 **de-configuring the computer system** to functionally remove the faulted
 14 component from the computer system as identified by said step of
 automatically identifying.

15 **1. “de-configuring the computer system”**

16 EMC asks the Court to construe this term as: “Arranging the computer system such that
 17 the faulted component is functionally removed from the system.” HP proposes (in its opposition
 18 brief): “Functionally removing a non-redundant faulted component from the computer system
 19 such that the function performed by the faulted component is not replaced. De-configuring the
 20 computer system does not include the step of replacing the faulted component or the functionality
 21 performed by the faulted component with a redundant component or with equivalent
 22 functionality.”

23 Nothing in the claim language indicates that only a “non-redundant” component or
 24 functionality is removed, and HP’s attempt to limit this term on the basis of the prosecution
 25 history is unsupported by the record. The inventors did not state that only “non-redundant”
 26 components are removed. Rather, in the cited portion of the prosecution history, they stated that
 27 in the claimed invention “a faulted component is not replaced by another redundant component

1 nor is the system re-configured.” Derringer Decl., Separate Appendix, vol. IX, p. 2611 (file
 2 history to U.S. Pat. No. 6,122,756).

3 The claim language indicates that one result of “de-configuring the computer system” is
 4 the removal of the identified faulted component from the computer system. Nothing in the claim
 5 language indicates that the component or its functionality is not replaced, although it is true that
 6 such replacement is not expressly claimed. Assuming for the sake of argument that the inventor
 7 disclaimed a “re-configuration” limitation in the course of prosecution, such an event would lead
 8 the Court only to omit such a limitation from protection by the claim at issue. It would not result
 9 in defining the “de-configuration” step by reference to a lack of “re-configuration.” The claim
 10 simply does not refer in any way to replacement of functionality or components.

11 Accordingly, the Court construes this term as “functionally removing a faulted
 12 component from the computer system.”

13 **2. “recovering”**

14 EMC argues that this term cannot be read as a limitation because it is in the preamble.
 15 Alternatively it asks the Court to construe the term as either: “recovering” or “restoring a system,
 16 program, database, or other system resource to a prior state following a failure or externally
 17 caused disaster.” HP proposes: “The act of resetting a system or data stored in a system to an
 18 operable state following damage, including rebooting the computer system, the operating system
 19 and application program(s).”

20 The preamble indicates that a computer system is automatically recovered after discovery
 21 of a faulted component. The body of the claim provides the steps for “recovering,” and, in so
 22 doing, it also indicates that a faulted component is automatically removed from a computer
 23 system.

24 It is well settled that if the body of the claim sets out the complete invention, and
 25 the preamble is not necessary to give life, meaning and vitality to the claim, then
 26 the preamble is of no significance to claim construction because it cannot be said
 to constitute or explain a claim limitation.

27 *Schumer v. Lab. Computer Sys., Inc.*, 308 F.3d 1304, 1310 (Fed. Cir. 2002). The preamble is
 28 redundant. The body of the claim provides that the system automatically discovers any faulted

1 components and removes them after doing so. Accordingly, the Court declines to construe this
 2 term.

3 **L. Disputed terms of the '347 patent.**

4 The text of claim 1 is set forth below with the key disputed terms highlighted in bold
 5 type.⁷

6 A system for automatically providing and maintaining data, said system
 7 comprising:

8 a host computer located in a first geographic location;

9 a first data storage system located in a first geographic location and coupled to
 10 said host computer, for storing data to be accessed by at least said host
 11 computer;

12 a second data storage system located in a second geographic location
 13 geographically remote from said first location, coupled to said first data
 14 storage system, for receiving at least data from said first data storage
 15 system; and

16 said first data storage system enabling transfer of said data to said second data
 17 storage system, concurrently with said data received from said host
 18 computer, so as to nearly simultaneously maintain a concurrent copy of
 19 data stored on said first data storage system and on said second data
 20 storage system wherein both said first and said second data storage
 21 systems **maintain an index**, said index including at least a first **indicator**
 22 providing an indication of whether a predetermined data element stored on
 23 said first data storage system is valid, and at least a second **indicator**
 24 providing an indication of whether said predetermined data element stored
 25 on said second data storage system is valid.

1 **1. "indicator"**

2 EMC asks the Court to construe this term as:

3 Information which provides a fairly certain sign or symptom of the validity of a
 4 predetermined data element. An indicator can include information about data
 5 located on a storage device, such as a track, or about data stored in an entire
 6 storage device via a device-level indicator. The requisite first and second
 7 indicators may be organized at different levels so long as the second provides
 8 information with respect to the validity of the same predetermined data element as
 9 the first even if that information is provided indirectly.

10 HP proposes: "An 'indicator' is something that 'indicates,' which means points out or points to,
 11 or states or expresses briefly, with certainty."

27
 28 ⁷ Although the Court refers to representative claim 1, it will construe the terms as used
 throughout the patent.

1 HP's proposal is an everyday definition of the word, except for the fact that it includes the
 2 requirement of "certainty." EMC proposes a construction that would limit the term to the
 3 "validity of a predetermined data element," but, instead of requiring "certainty," would require
 4 only an "indication" that is "fairly certain." The claim language states that an "indicator"
 5 provides an indication of whether a predetermined data element stored on a first or second data
 6 storage system is valid, thus supporting EMC's proposed "predetermined data element"
 7 limitation. However, that limitation stands on its own and need not be included in the Court's
 8 specific construction of the term "indicator."

9 The claim language does not help to resolve the parties' dispute over the level of certainty
 10 required. Indeed, nothing in the claim language explicitly limits the term to "indicates with
 11 certainty." The word "indicate," by itself, does not necessarily require certainty; in fact, HP's
 12 addition of modifying words suggests that it seeks a particularly narrow construction. The
 13 dictionary definition of "indicate" does not include "certainty." ("Indicate" means "to point out
 14 or point to," or "to be a sign, symptom, or index of." Webster's Ninth New Collegiate
 15 Dictionary.)

16 Nor is there a requirement of certainty in the specification, which uses the term
 17 "indicator" according to its common meaning. For example, the "secondary data storage system
 18 controller provides an indication or acknowledgement [sic] to the primary data storage system
 19 controller that the primary data to be copied to the secondary data storage system in identical
 20 form as secondary data has been received or . . . written." '347:2/63-3/1. The invention is able
 21 to "achieve[] nearly 100 percent data integrity." '347:2/34-35. The "indicators" also indicate
 22 several measures of "validity." There are indications of "write or copy pending . . . of both the
 23 primary data . . . and the secondary data," '347:7/46-47, "pending format change," '347:7/48, and
 24 the validity of data storage location, '347:8/10-12. There are also "indicators" of whether a
 25 primary or secondary device may be written to and whether a primary and secondary track is not
 26 shown. '347:8/23-24 & 8/36. In none of these discussions of the various indicators is there a
 27 suggestion or requirement that the indicator indicate with certainty. Such a requirement likely
 28 would be a highly significant limitation and improvement, as very few (and possibly no)

1 computer systems ever function with absolute certainty.

2 Accordingly, the Court construes “indicator” as “something that ‘indicates,’ which means
3 points out or points to, or to be a sign, symptom, or index of.”

4 **2. “maintain an index”**

5 EMC argues that this term should be given its plain and ordinary meaning. Alternatively
6 it asks the Court to construe it as: “Keep or store an index.” HP proposes: “To keep an index in
7 an existing, up-to-date state (as of repair, efficiency or validity).” HP thus seeks to include a
8 limitation requiring the index to be kept in an “up-to-date state.”

9 It is unclear what “up-to-date-state” would mean in the context of the claim. It is true that
10 the claim indicates that the index is modified with use. For example, the claim language suggests
11 that the index is updated as the storage systems store new data: the

12 storage systems maintain an index, said index including at least a first indicator
13 providing an indication of whether a predetermined data element stored on said
14 first data storage system is valid, and at least a second indicator providing an
indication of whether said predetermined data element stored on said second data
storage system is valid.

15 However, the claim language does not specify *when* the index is updated.

16 Similarly, although several portions of the written description suggest that the index is
17 updated, the specification provides no basis for the “up-to-date-state” limitation, nor does it
18 provide guidance as to what such language would mean. For example, the storage system
19 controller “maintaining a list of primary data to be copied updates this list to reflect that the given
20 primary data has been received by and/or copied to the secondary data storage system.”

21 ‘347:3/20-23.

22 [S]ince data will not be immediately synchronized between the primary and
23 secondary data storage systems, data integrity must be maintained by maintaining
24 an index or list of various criteria including a list of data which has not been
mirrored or copied, data storage locations for which a reformat operation is
25 pending, a list of invalid data storage device locations or tracks, whether a given
device is ready, or whether a device is write-disabled. Information must also be
included as to the time of the last operation so that the data may later be
26 synchronized should an error be detected.

27 ‘347:7/22-32. Moreover, “[w]hen the primary data storage system controller disk adapter writes
28 the data to the primary data storage device, it will reset bit 104 of the write pending indicator

1 bits.” ‘347:7/59-62. Thus, although the specification describes a system in which indices are
2 updated as part of the data storage on both systems, it does not include a requirement as to when
3 such updating must or should occur. Under these circumstances, importation of a limitation
4 requiring the index to be kept in an “up-to-date-state” could exclude the claimed invention itself
5 if the invention’s indices were not updated *instantaneously*.

6 In the absence of contrary language in the specification, a person having ordinary skill in
7 the art would understand this term to have its plain meaning. A dictionary definition—consistent
8 with the description of the claimed invention—of “maintain” is “to keep in an existing state (as
9 of repair, efficiency, or validity).” Webster’s Ninth New Collegiate Dictionary. According to the
10 plain meaning of the term, the index must be kept in a state of proper use as defined by the claim.
11 The claim indicates that the index contains information concerning the validity of predetermined
12 data elements. Thus, because the index generally must represent the existing states of the storage
13 systems, it must be updated at some point after the validity of predetermined storage elements is
14 altered.

15 Accordingly, the Court construes this term as “keeping—although not necessarily
16 instantaneously—an index in a state of validity.”

17
18 IT IS SO ORDERED.

19
20 DATED: June 23, 2004

21 /s/ (electronic signature authorized)
22 JEREMY FOGEL
23 United States District Judge
24
25
26
27
28

1 Copies of Order sent to:

2 Stephen N. Adams
sadams@orrick.com
3 marlantico@orrick.com
aako-nai@orrick.com

4 Todd Michael Briggs
5 tbriggs@orrick.com

6 I. Neel Chatterjee
nchatterjee@orrick.com
7 kmudurian@orrick.com

8 Mark E Ferguson
mark.ferguson@bartlit-beck.com

9 Robert T. Haslam
10 rhaslam@hewm.com

11 Chris Landgraff
chris.landgraff@bartlit-beck.com

12 Chris R. Ottenweller
cottenweller@orrick.com
13 hbursik@orrick.com

14 Michael K. Plimack
15 mplimack@hewm.com

16 Amy Kathleen Van Zant
avanzant@hewm.com

17 Lisa C. Ward
18 lward@orrick.com

19 Mark R. Weinstein
mweinstein@orrick.com

20

21

22

23

24

25

26

27

28